THE UNIVERSITY OF CHICAGO

EXPANDERS WITH SYMMETRY: CONSTRUCTIONS AND APPLICATIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

BY

TUSHANT MITTAL

CHICAGO, ILLINOIS
AUGUST 2024

To people who loved me too much and left me too soon:
my grandmother, Taramani, and my dear friend, Shreyas.

TABLE OF CONTENTS

LIST OF TABLES

ACKNOWLEDGMENTS

I am indebted to numerous people for making the last six years of my life delightful and helping me grow as a researcher and a person.

I have received nothing but constant encouragement from my advisors, Janos Simon, and Madhur Tulsiani. Janos has imparted wisdom to me in his characteristically cheerful manner since I met him during visit day. Madhur has been a caring advisor, a wise mentor, and, most importantly, a person I could look up to in aspects beyond research. His friendly attitude, readiness to help, and ability to give tailored advice in the nicest way possible have always impressed me. Half the enjoyment of my grad life was talking to Madhur about everything under the sun, an activity that invariably made me feel better. I hope to imbibe all he has taught and carry it with me forever.

I want to thank my committee members, Bill Fefferman (also on my Master's committee) and Ryan O'Donnell (also a collaborator), for contributing their time, providing valuable feedback, and pointing me to interesting future directions. I am grateful to the National Science Foundation (NSF) for supporting my research (NSF grant CCF-1816372) and the Simons Institute for the Theory of Computing for hosting me more than once.

I have had the privilege of working with various people who have imparted invaluable lessons. For example, it was enlightening to witness Gabor get to the core of a proof via an elegant definition. An exhaustive list of my learnings would be too long, so here is a chronological list of my wonderful collaborators and teachers: Kapil Paranjape, Nitin Saxena, Matilde Lalín, Ketan Mulmuley, Gabor Ivanyos, Youming Qiao, Pallav Goyal, Rafael Oliveira, Abhibhav Garg, Madhur Tulsiani, Fernando Granha Jeronimo, Pedro Paredes, Ryan O'Donnell, Sourya Roy, Avi Wigderson, Shashank Srivastava, Max Hopkins, and Suprovat Ghoshal.

Having advisors at two institutions comes with its perks, not the least of which is having twice the number of helpful administrators to assist you! Nita, Megan, Donna,

and Margaret at the University of Chicago and Mary and Amy at TTIC have promptly addressed my queries, enabling me to focus on research. I also got to be a part of two excellent departments. I will always cherish the grad lounge hangouts at UChicago with Alex, Anne, Ashka, Bogdan, Casey, Kavon, Francesca, Jean, Jonathan, Kartik, Ryan, Suhail, and Valerie. The dimly lit "theory room" was, of course, my usual haunt, shared with other fantastic theorists: Alex, Antares, Aritra, Chris, Erasmo, Gabe, Goutam, Jafar, Jesse, Joseph, Konstantinos, Kunal, Matt, Nathan, Neng, Owen, Roozbeh, and Theo. Subjecting Alex, Chris, and Neng gleefully to my “hot takes” was a favorite hobby of mine. The theory group at TTIC adopted me and let me into their inner circle, comprising Anmol, David, Kavya, Kshitij, Naren, Max, Pushkar, Saba, Shashank, and Sudarshan. I have enjoyed hanging out with this gang at daily tea-time sessions (special thanks to Kavya for the fancy teas), impromptu dinners, and festive celebrations.

Most of my time outside of work has been spent with a motley crew of friends who have not just entertained my shenanigans but also constantly pushed me to do better. In no particular order: Aravind, a companion since my undergrad days through whom I have read a lot of self-help books; Kshitij, my dearest roommate of five years who has relentlessly tried to infect me with his optimism; Pallav, the dependable friend whose presence always soothed me in moments of crisis, both mathematical and personal; Amruta, the social butterfly who has constantly cheered me on while teaching me essential skills; and Jafar who showered me with a unique dose of brotherly love, good-natured ribbing, and sagely advice. I have additionally received such warmth and hospitality from my more recent friends—Ahona, Amrita, Darpan, Ishika, Marium, Sourya, Varsha, and Zara—that graduating has become more bittersweet than I had imagined!

Moving on to familial bonds, my cousin Palak has been my guide to the American way of life and a steadfast pillar of support from the moment I arrived in Chicago. Her apartment has been an oasis of comfort and invaluable advice through all the turmoil of

grad life, and I will forever miss our late-night chats.

I am incredibly grateful to my parents for everything they have given me. My dad has been an anchor of strength who has always had my back, especially during the lowest points of grad school. Socratic dialogue-styled phone calls with my mom, the wisest person I know, have been essential for me to make sense of the world around me. My brother, Kaushal, always manages to lighten up my mood, and I cannot thank him enough for shielding me from the turbulence of family life and letting me stay aloof. My extended family, especially my aunt Usha, has showered me with much affection, for which I am immensely thankful.

Chicago has been an ideal host with its majestic lake, vibrant cultural life, and peaceful winters that create the perfect conditions to ponder. To ponder upon the beauty of mathematics is what I do research for, and I am grateful to my teachers and the lovely yellow books that have given me a taste of this beauty. Another perennial source of beauty in my life has been poetry, which is quite like math in that they both generate insight aided by a mysterious process of symbol shifting. I owe a great debt of gratitude to the poets—especially Ghalib and Iqbal—whose verses have taken a seat as a voice in my head, giving me company and assisting me with the lonely, everyday task of meaning-making.

ABSTRACT

Expanders are sparse yet well-connected graphs with numerous theoretical and practical uses. Symmetry is a valuable structure for expanders as it enables efficient algorithms and a richer set of applications. This thesis studies expanders with symmetry, giving new constructions and applications.

We extend expander construction techniques to work with symmetry and give explicit constructions of expanders with varying quality of expansion and symmetries of various groups. In particular, we construct graphs with large Abelian group symmetries via the technique of *graph lifts*. We also give a generic amplification procedure that converts a weak expander to an almost optimal one while preserving symmetries. This procedure is obtained by generalizing prior amplification techniques that work for Cayley graphs over Abelian groups to Cayley graphs over any finite group. In particular, we obtain almost-Ramanujan expanders over every non-abelian finite simple group.

We then explore the utility of having both symmetry and expansion simultaneously. We obtain explicit quantum LDPC codes of almost linear distance and *good* classical quasi-cyclic codes with varying circulant sizes using prior results and our constructions of graphs with Abelian symmetries. We show how our generic amplification machinery boosts various structured expander-like objects: *quantum expanders*, *dimension expanders*, and *monotone expanders*. Finally, we prove a structural result about expanding Cayley graphs, showing that they satisfy a "degree-2" variant of the *expander mixing lemma*. As an application of this, we give a randomness-efficient query algorithm for *homomorphism testing* of unitary-valued functions on finite groups and a derandomized version of the celebrated Babai–Nikolov–Pyber (BNP) lemma.

CHAPTER 1

INTRODUCTION

> *One of my personal beliefs is that fascination with symmetries and groups is one way of coping with frustrations of life's limitations: we like to recognize symmetries which allow us to recognize more than what we can see.*
>
> Pierre de la Harpe, Topics in Geometric Group Theory

Graphs are structures that capture pairwise relations between elements of a set, and many real-world tasks can be modeled as a graph-theoretic problem. Computer scientists are often interested in developing primitives such as error correction algorithms or cryptographic schemes. A successful approach to many of these constructions requires designing graphs with special properties.

One such property is that of *expansion*, a major development in theoretical computer science over the past few decades that has led to numerous algorithmic advances, lower bounds in complexity theory, and a lot more (see [HLW06] for a comprehensive survey). Expansion is the magical property of being sparse yet very well-connected, and the quest to construct expander graphs has led to structural results in pure mathematics [Lub12].

Symmetry is often a desirable property for any object, and the study of graphs with symmetry can be motivated by at least two different streams of research. One is the topic of *geometric group theory*, an approach to understanding groups via the spaces it acts on, such as graphs. From a computer science perspective, one aims to utilize the symmetry of graphs to improve graph-based constructions and algorithms. This thesis concerns graphs with both expansion and symmetry, and addresses the following questions:

How can we explicitly construct expanders with given symmetries?

How can one utilize expanders with symmetries?

These two properties have a synergistic relation, which makes studying them in conjunction with each other fruitful. For instance, all known constructions of optimal (non-bipartite) expanders (*Ramanujan expanders*) are highly symmetric constructed via group theory. Moreover, symmetry enriches graph-based constructions, even when not necessary. For example, error-correcting codes can be built from expanders without any requirements on symmetry, and such codes have great error-correction capacity. However, if one requires the code to be *locally testable*—a property that enables very efficient detection of errors—the only known way is via symmetry.

In the other direction, the study of symmetry, i.e., group theory, is also enriched by the perspective of expansion. A *Cayley graph*, $\mathrm{Cay}(G, S)$, is a group-theoretic graph defined using a group G and a subset $S \subseteq G$. This construction provides a link between graphs and groups, and one can now ask questions like, which groups admit expanding Cayley graphs? Studying such questions has not only led to expander constructions but has also shed light on the properties of these groups. For instance, answering this question for one class of groups[1] involved developing structural results in group theory like *product theorems*, variants of *Kazhdan's property (T)*, the notion of *quasirandomness*, etc.

The second part of this thesis gives a few more concrete applications of such graphs. Before we march on, let us formalize these terms. We will always work with an infinite family of d-regular undirected graphs $\{X_i\}_{i\in\mathbb{N}}$, where d is an absolute constant and the size of the graphs grows with i. This ensures that the graphs are very sparse.

We will work with the *spectral* notion of expansion. For a graph X, $\lambda(X)$ is the second largest singular value of its normalized adjacency matrix, A_X. We say that a family of d-regular graphs, $\{X_i\}_{i\in\mathbb{N}}$, is a (d, λ)-expander (or just λ-expander), if $\lambda(X_i) \leq \lambda$ for every member X_i of the family. The smaller the expansion parameter λ, the more spectrally expanding the family. The trivial bound is $\lambda = 1$, and we say a graph is an expander

1. The class of non-abelian finite simple groups.

if $\lambda < 1$. This quality of expansion suffices for most applications. When working with graphs where we need to analyze the dependence of λ on d, it can be more convenient to work with the unnormalized eigenvalue, $\lambda_u := \lambda \cdot d$.

To formalize symmetry, first recall that a graph isomorphism is a permutation of the vertices that preserves the graph structure. Namely, a bijection $\sigma : V \to V$ such that (u, v) is an edge if and only if $(\sigma(u), \sigma(v))$ is. Let $\mathrm{Aut}(X)$ be the group of all graph isomorphisms of the graph X. We say that X has symmetries of G if $G \subseteq \mathrm{Aut}(X)$. In other words, G acts on the graph, i.e., for each $g \in G$, we have a graph isomorphism $\varphi(g)$ of X such that $\varphi(gh) = \varphi(g)\varphi(h)$. The formal problem statement now is the following:

Question 1.0.1 (Expanders with Symmetry)**.** *Fix a family of groups* $\{G_i\}_i$*, a positive integer* d*, and a real number* $\lambda \in [0, 1)$*. Give a deterministic polynomial-time algorithm that for each* $i \in \mathbb{N}$*, computes a* d*-regular graph,* X_i *such that* $\lambda(X_i) \leq \lambda$*, and* $G_i \subseteq \mathrm{Aut}(X_i)$*.*

A prominent application of expanders is in the area of *pseudorandomness*. Algorithms often use randomness, a precious resource, and hence, the desire to reduce the amount of randomness. Say that the algorithm computes a function $f : G \to \mathbb{C}$ by sampling random inputs. We say that f is δ-fooled by a set $S \subseteq G$ if,

$$\left| \mathop{\mathbb{E}}_{x \sim G}[f(x)] - \mathop{\mathbb{E}}_{x \sim S}[f(x)] \right| \leq \delta.$$

Thus, the function f can be approximated (in expectation) by the uniform distribution over S instead of the uniform distribution over G. Sampling from S requires $\log(|S|)$ bits of randomness, which can be much smaller than $\log(|G|)$ required to sample from G.

This question about groups, yet again, relates to graphs: $\mathrm{Cay}(G, S)$ is a λ-expander if and only if S λ-fools all *representations* of G. This equivalence poses an interesting question:

Question 1.0.2 (Pseudorandomness of Symmetric Expanders)**.** *Let* G *be a finite group and* $S \subseteq G$ *such that* $\mathrm{Cay}(G, S)$ *is a* λ*-expander. Determine the set of functions,* $f : G \to \mathbb{C}$ *that are* $\delta(\lambda)$*-fooled by* S*, i.e.,* $\left| \mathbb{E}_{x \sim G}[f(x)] - \mathbb{E}_{x \sim S}[f(x)] \right| \leq \delta(\lambda)$*.*

1.1 Technical overview

We now present an impressionistic sketch of a few essential concepts and techniques this thesis deals with.

Cayley Graphs A natural way to have symmetry of a group G is to have the graph's vertex set be G. The action then is merely by (left) group multiplication. Since this action is a group isomorphism, we get that (g, h) is an edge if and only if $(1, g^{-1}h)$ is. Thus, the graph is determined by the edges going out of the identity element, 1. This construction is called the Cayley graph, denoted as $\mathrm{Cay}(G, S)$. Here, $S \subseteq G$ is a symmetric multiset, denoting the neighbors of 1. Thus, $g, h \in G$ are adjacent if and only if $g^{-1}h$ belongs to S. The set S is called the *generating set*, as S generates G as a group.

Group-based Graph Lifts One can generalize the above construction by having $V = V_0 \times G$, where the action is merely on G. In this setup, one has many more choices. In particular, one can first pick $E_0 \subseteq V_0 \times V_0$ arbitrarily, and then for each $(u, v) \in E_0$, pick a subset $S_{u,v} \subseteq G$. Then, the edges are $E = \{((u, g), (v, h)) \mid (u, v) \in E_0,\ g^{-1}h \in S_{u,v}\}$. As we wish to vary the family of groups $\{G_i\}$, this graph can be seen as the *lift* of a base graph $X_0 = (V_0, E_0)$. Moreover, we wish to preserve the degree, and so, $S_{u,v}$ is a singleton[2]. Thus, we can generate $\{X_i\}_i$ where $V_i = V_0 \times G_i$ and the edges are defined using the *signing function* $s : E_0 \to G_i$ that maps $(u, v) \mapsto \{s(u, v)\} = S_{u,v}$.

Bootstrapping Expansion In the above constructions, the graph depends on the choice of the generating set S, or the *signing function*. It is pretty challenging to build a generating set from scratch (even randomly) that yields an expander, and most techniques use deep results from number theory and representation theory [LPS88, Mar88, Mor94]. Com-

2. One only needs $S_{u,v}$ to be constant-sized. However, requiring it to be a singleton is not without loss of generality as we can replace the edge (u, v) with multiple copies, which has the same effect.

binatorial techniques [RVW00, BL06, BATS08, MOP20, OW20, Alo21] start with a weak expander and amplify it—either in size or quality—which is significantly easier. However, these amplification operations do not preserve any algebraic structure, as they are focused only on expansion. One key contribution of this thesis is to utilize these combinatorial techniques to amplify expansion *while preserving algebraic structure*, i.e., symmetry. For example, to improve the expansion of a Cayley graph $\mathrm{Cay}(G, S)$, we are not allowed to modify the graph arbitrarily but only to change S. This preserves the Cayley structure and, thus, the symmetries of G.

Spectral norm and its non-abelian analog The ℓ_1-norm of the Fourier transform of a function is known as its *spectral norm*. Spectral norm has emerged as an important quantity for the analysis of Boolean functions, i.e., functions over $\mathbb{Z}_2^n$. In particular, functions with low spectral norm have a lot of structure [STV17]: they admit small decision trees, parity decision trees, they are easily learnable, etc. One of the contributions of this thesis is to study the non-abelian analog of this norm from the perspective of pseudorandomness. A first generalization one can think of would be a similar ℓ_1 norm of the Fourier coefficients. However, it turns out that the appropriate generalization of the spectral norm is the *Fourier algebra norm*. This was suggested earlier by Sanders [San21], who used it to generalize the quantitative idempotent theorem. This norm has multiple equivalent definitions, but our key idea is to use the following harmonic analytic reformulation due to Eymard [Eym64],

$$\|f\|_A = \min_{(\pi,V)} \left\{ \|u\| \cdot \|v\| \;\middle|\; f(x) = \langle u, \pi(x)\, v \rangle \right\},$$

where $u, v \in V$, and (π, V) is a representation of G, i.e., $\pi : G \to \mathbb{U}_V$ is a homomorphism from the group G to the group of unitary operators on the complex Hilbert space V.

It is well-known that any function, f, on an Abelian group is $\varepsilon\|\hat{f}\|_1$-fooled by any ε-biased set. We show that this neatly generalizes to any finite group via the Fourier algebra norm, i.e., any function, f, on a finite group is $\varepsilon\|f\|_A$-fooled by any ε-biased set.

1.2 Key Contributions

1.2.1 Expanders with Abelian Symmetries via Lifts

Using group-based graph lifts, we construct expanders with large Abelian symmetries. For this result, we use the eigenvalue of the unnormalized adjacency matrix, $\lambda_u(X)$.

Theorem 1.2.1 (Explicit Abelian Lifts)**.** *For large enough* n *and constant degree* $d \geq 3$*, given an abelian group* G*, and any fixed constant* $\varepsilon \in (0,1)$*, we can construct a* d*-regular graph* X *on* $\Theta(n|G|)$ *vertices, in deterministic polynomial time, such that,*

1. X *is* G*-lift of a graph* X_0 *on* $\Theta(n)$ *vertices. Thus,* $G \subseteq \mathrm{Aut}(X)$*.*
2. *If* $|G| \leq \exp\left(n^{\delta(d,\varepsilon)}\right)$*, then* $\lambda_u(X) \leq 2\sqrt{d-1} + \varepsilon$*.*
3. *If* $|G| \leq \exp\left(n^{\delta}\right)$ *and also* $d \geq d_0(\varepsilon)$*, then* $\lambda_u(X) \leq \varepsilon \cdot d$*.*
4. *If* $|G| \leq \exp\left(cnd^{-\frac{1}{2}}\right)$*, then* $\lambda_u(X) \leq O(\sqrt{d}\log d)$*.*
5. *If* $|G| = \exp\left(cnd^{\delta}\right)$ *for* $\delta \in [-1/2, 1)$*, then* $\lambda_u(X) \leq O\left(d^{\frac{2+\delta}{3}}\log d\right)$*.*

This result is proved in two parts using different proof techniques. In Chapter 2, we give a simple proof via *discrepancy* by utilizing the work of Bilu and Linial [BL06]. This gives an algorithm that runs in time $\mathrm{poly}(2^n, |G|)$ which is efficient when $|G| = 2^{\Omega(n)}$. To tackle smaller lifts, we use the *trace method* in Chapter 3, generalizing the results of [MOP20] for $G = \mathbb{Z}_2$. This gives the above result for the sub-exponential sizes (regimes 2 and 3 above).

1.2.2 Expanders with non-Abelian Symmetries via Amplification

Any infinite family of d-regular expanders satisfies, $2\sqrt{d-1} - o(1) \leq \lambda_u(X)$, the Alon-Boppana bound [Nil91]. Graphs that achieve this bound are called Ramanujan graphs, which are optimal expanders in this spectral sense. We will use the normalized eigenvalue and refer to the weaker bound of $d \leq O\left(\lambda^{\frac{-1}{2}+o(1)}\right)$, as *almost Ramanujan*.

We give a generic amplification procedure that converts a weak expander to an almost Ramanujan one while preserving symmetries. The key technical result is to establish such an amplification for Cayley graphs over arbitrary finite groups,

Theorem 1.2.2 (Amplifying Cayley Graphs)**.** *Let* G *be a finite group and* S *be such that* $\mathrm{Cay}(G, S)$ *is a* λ_0*-expander, for some constant* $\lambda_0 \in (0, 1)$*. For every* $\lambda \in (0, 1)$*, there exists* S' *such that,*

- $\mathrm{Cay}(G, S')$ *is a* λ*-expander.*
- $|S'| = O\left(|S|/\lambda^{2+o_\lambda(1)}\right)$*, and*
- S' *can be computed deterministically in* $\mathrm{poly}(|S|/\lambda)$*-time assuming an oracle for group operations.*

Furthermore, if $\mathrm{Cay}(G, S)$ *is strongly explicit*[3]*, then so is* $\mathrm{Cay}(G, S')$*.*

Breuillard and Lubotzky [BL22] ask whether having near-Ramanujan expanders for all families of non-abelian finite simple groups is possible. Theorem 1.2.2 makes progress towards this question (the $o(1)$ term needs to be removed to resolve it completely). Interestingly, the above result for Cayley graphs implies an expansion result for general families of (regular) expander graphs. The key idea is to use a result of König that says that the adjacency matrix, say A_X, of an arbitrary regular graph, can be written as a sum of permutation matrices. This gives the following,

Theorem 1.2.3 (Amplifying General Expanders)**.** *Let* $\{X_i\}_{i\in\mathbb{N}}$ *be a family of* (d_0, λ_0)*-expanders where* $\lambda_0 < 1$ *is a constant. For any (target)* $\lambda \in (0, 1)$ *and* X_i*, we can explicitly*[4] *construct a* (d, λ)*-expander,* X'_i*, on the same vertex set, where* $d = O(d_0/\lambda^{2+o_\lambda(1)})$*. Moreover, the construction is local in the sense that edges in* X'_i *correspond to short walks in* X_i*.*

3. Neighbors of a vertex can be computed in time polynomial in the *description length* of a vertex.

4. See Definition 1.3.6

1.2.3 *Applications of Expanders with Symmetries*

The second part of the thesis explores how one can use these symmetric expanders to design randomness-efficient algorithms and construct other interesting pseudorandom objects. We do not state the theorems formally here, as they require many new definitions.

In Chapter 5, we look at useful corollaries of our symmetric expander constructions. We start with the original motivation to embark on this topic, constructing large-distance quantum error-correcting codes. We will see how our construction of graphs with abelian symmetry (Theorem 1.2.1) plugs into the machinery of Kalachev and Panteleev [PK21] to yield explicit quantum codes with almost linear distance. We then explore the amplification of various structured expander-like objects, like *quantum expanders*, via our amplification of Cayley graphs (Theorem 1.2.2).

In Chapter 6, we address Question 1.0.2 and show that a Cayley expander over G can fool functions $f : G \to \mathbb{C}^{t\times t}$ with small algebra norm. This leads to a randomness-efficient query algorithm for testing if a function to $t \times t$ unitary matrices, $f : G \to \mathbb{U}_t$, is a *homomorphism*. Prior algorithms sample a uniformly random pair (x, y) and test if $f(xy) = f(x)f(y)$, which is the homomorphism property. This uses $2\log|G|$ bits of randomness, and our contribution is to reduce this by showing that one can instead sample a pair using edges of a Cayley expander over G. This reduces the randomness to $\log|G| + \log|S|$ which is $(1 + o(1))\log|G|$, i.e., almost optimal.

The main technical contribution is to show that expanding Cayley graphs satisfy a "degree-2" variant of the *expander mixing lemma* (EML). This degree-2 EML also gives a derandomized version of the Babai–Nikolov–Pyber (BNP) lemma.

1.3 Preliminaries

Operators

We will mostly work with operators over finite-dimensional Hilbert spaces, i.e., vector spaces with an inner product. However, in Section 5.2.6, we will work with infinite-dimensional spaces, and thus, we provide general definitions.

Definition 1.3.1 (Spectrum)**.** For an operator T, define, $\mathrm{Spec}(T) = \{\lambda \mid (T - \lambda I) \text{ is invertible}\}$. We define the *spectral radius* $\rho(T) = \max_{\lambda \in \mathrm{Spec}(T)} |\lambda|$. When the spectrum is discrete, i.e., $\mathrm{Spec}(T) = \{|\lambda_1(T)| \geq \cdots \geq |\lambda_n(T)|\}$, define $\rho_2(T) = |\lambda_2(T)|$.

Definition 1.3.2 (Operator Norm)**.** Let $\mathcal{H}, \mathcal{H}'$ be Hilbert spaces with norms induced from their respective inner products. For any bounded linear operator $T : \mathcal{H} \to \mathcal{H}'$, we define,

$$\|T\|_{\mathrm{op}} = \sup_{v \in \mathcal{H}} \frac{\|Tv\|}{\|v\|} = \sup_{v \in \mathcal{H}\, w \in \mathcal{H}'} \frac{|\langle Tv, w\rangle|}{\|v\|\|w\|}.$$

Alternatively, $\|T\|_{\mathrm{op}}^2 = \rho\left(TT^*\right)$ where T^* is the adjoint of T.

Graphs and Spectral Expanders

Throughout this thesis, we will use $X = (V, E)$ to denote an n-vertex d-regular undirected multigraph for some $d \geq 1$. We denote by A_X the normalized adjacency matrix of X.

Definition 1.3.3 (λ-spectral Expander)**.** We say that X is a λ-spectral expander if $\rho_2(A_X) \leq \lambda$. We will use the notation $\lambda_u(X) = \rho_2(A'_X)$ where A'_X is the unnormalized adjacency matrix.

Lemma 1.3.4 (Expander Mixing Lemma)**.** *Let S, T be subsets of the vertices of a d-regular graph. Define $E(S, T) = \{(x, y) \in E \mid x \in S,\ y \in T\}$. Then,*

$$\left||E(S,T)| - \frac{d|S||T|}{n}\right| \leq \lambda_u(G)\sqrt{|S||T|}.$$

Definition 1.3.5 (Cayley Graph)**.** Let G be a finite group and $S \subseteq G$ be a generating set such that if $s \in S$, then $s^{-1} \in S$. Then, $\mathrm{Cay}(G, S)$ is an undirected graph with vertex set G, and $(g, h) \in E$ if and only if $gh^{-1} \in S$.

Definition 1.3.6 (Explicit graph)**.** A family of graphs $\{X_i\}_{i \in \mathbb{N}}$ is said to be *explicit* if the adjacency matrix of X_i can be computed deterministically in $\mathrm{poly}(|X_i|)$-time. Moreover, it is said to be *strongly explicit* if the list of its neighbors of any vertex in X_i can be computed $\mathrm{poly}(\log |X_i|)$-time.

Group Representations and Small-Bias sets

For finite groups, every representation can be made unitary; thus, studying these suffices. Let V be a complex Hilbert space and denote by U_V, the unitary group of operators acting on V.

Definition 1.3.7 (Unitary Group Representation)**.** For a group G, a unitary representation is a pair (ρ, V) where $\rho : G \to U_V$ is a group homomorphism, i.e., for every $x, y \in G$, we have $\rho(xy) = \rho(x)\rho(y)$. A representation is *irreducible* if the only subspaces of V that are invariant under the action of $\rho(G)$ are the empty space, $\{0\}$, and the entire space, V.

For a representation (ρ, V), will use d_ρ to denote $\dim(V)$. We use $\widehat{G}$ to denote the set of all irreducible representations (*irreps*) of a group G. Every group has two special irreducible representations:

- The *trivial representation*, $(\rho_{\mathrm{triv}}, \mathbb{C})$, where $\rho(g) = 1$ for every group element $g \in G$.
- The *regular representation*, $(\rho_{\mathrm{reg}}, \mathbb{C}[G])$ where $\rho(g)\, e_h = e_{gh}$ for every $g, h \in G$.

The following is a fundamental result that states that every representation decomposes as a finite sum of irreducible ones.

Theorem 1.3.8 (Maschke)**.** *Let* G *be a finite group and let* (π, V) *be any representation of* G*. Then,* $V = \oplus_i V_{\rho_i}$*, i.e., it decomposes as a direct sum of irreducible representations* $\{\rho_i\}_i$*. Explicitly, there exists a unitary transformation* U_π *such that* $U_\pi \pi U_\pi^*$ *is block-diagonal with each block being* ρ_i*.*

Fact 1.3.9 (Decomposing the Regular Representation)**.** *For any finite group* G*, the regular representation contains all irreps, with multiplicity exactly* d_ρ*, i.e.,* $\mathbb{C}[G] = \oplus_{\rho \in \widehat{G}} V_\rho^{d_\rho}$*.*

Small–bias sets and Pseudorandomness For a group G, small-bias sets are multisets (or distributions), $S \subseteq G$ that fool all non-trivial irreducible representations. Small bias sets over Abelian groups were introduced in the pioneering work of Naor and Naor [NN93] and are a fundamental derandomization tool widely used across domains like complexity theory, coding theory, learning theory, graph theory, etc.

Definition 1.3.10 (ε-Biased Set)**.** Let $\varepsilon \in [0, 1)$. We say that a multiset $S \subseteq G$ is ε-biased if for every irreducible representation ρ, ρ is ε-fooled by S, i.e., $\|\mathbb{E}_{s \sim S}[\rho(s)]\|_{\mathrm{op}} \leq \varepsilon$.

Fact 1.3.11. *A multiset* $S \subseteq G$ *is an* ε*-biased set if and only if the* $\mathrm{Cay}(G, S)$ *is an* ε*-expander.*

Proof. The normalized adjacency matrix is $A_\chi = \mathbb{E}_s \rho_{\mathrm{reg}}(s)$. Now apply Theorem 1.3.8 and Fact 1.3.9. □

1.4 Bibliographic Note

All the results in this thesis are joint works.

1. Chapters 2 and 3, Section 5.1, and Appendices A.2 and A.3 are based on the work: Fernando Granha Jeronimo, Tushant Mittal, Ryan O'Donnell, Pedro Paredes, and Madhur Tulsiani. Explicit Abelian Lifts and Quantum LDPC Codes. In *13th Innovations in Theoretical Computer Science Conference* (ITCS 2022), 2022.
   DOI: 10.4230/LIPIcs.ITCS.2022.88

2. Chapter 4, Section 5.2, and Appendix A.1 are based on the work:

   Fernando Granha Jeronimo, Tushant Mittal, Sourya Roy, and Avi Wigderson. Almost Ramanujan Expanders from Arbitrary Expanders via Operator Amplification. In *Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science*, 2022. DOI: 10.1109/FOCS54457.2022.00043

3. Chapter 6 is based on the work:

   Tushant Mittal and Sourya Roy. Derandomized Non-Abelian Homomorphism Testing in Low Soundness Regime, 2024. Preprint, arXiv : 2405. 18998 [cs . CC]. DOI: 10.48550/arXiv.2405.18998

Part I

Constructions

CHAPTER 2

GRAPH LIFTS VIA DISCREPANCY

Laqa: Question for you, Piro. I give you a sparse binary matrix, A, and randomly flip each 1 to a -1. Now, give me a bound on the eigenvalue of this signed matrix.

Piro: Okay, so it is a Bernoulli sum of matrices. What if I apply Matrix Chernoff?

Laqa: That is good but not good enough! You get bounds depending on the dimension of the matrix, but I need dimension-independent bounds.

Piro: For dimension independence I need more information about A. What is the context?

Laqa: A is the adjacency matrix of an expander, and this random matrix is what you get when you analyze a random 2 lift.

Piro: In that case, the expander mixing lemma tells you that there are many edges across a cut, so you can get Chernoff-like concentration for $u^\intercal A v$ where u, v denote indicators of sets. And then, try to take a union bound.

Laqa: Interesting, but I need a bound on this Rayleigh quotient for all real vectors.

Piro: That is what ε-nets are for! Take a look at the original paper of Bilu and Linial.

2.1 Introduction to Graph Lifts

Let $X = (V, E)$ be a graph and assume that we have an ordering on V and, by convention, $(u, v) \in E$ if $u \leq v$. Let E^d denote the set of directed edges i.e. $E^d = \bigcup_{(u,v)\in E}\{(u, v), (v, u)\}$. For a group G, a G-signing is a function $s : E^d \to G$ such that $s(v, u) = s(u, v)^{-1}$.

Definition 2.1.1 (G-lift of a graph)**.** Given a G-signing of an undirected graph $X = (V, E)$, the lifted graph $X(s) = (V', E')$ is a graph on $|G|$ copies of the vertices $V' = V \times G$ where for every edge $(u, v) \in E^d$ we have $((u, i), (v, s(u, v)\cdot i)) \in E'$.

A lift is called *random* if the signing function is defined at random, i.e., for each $(u, v) \in E$, we assign $s(u, v)$ uniformly at random from G. The following matrix is crucial in the analysis of lifted graphs.

Definition 2.1.2 (Signed matrix)**.** Let A be the adjacency matrix of a graph X. For a character $\chi : G \to \mathbb{C}$, define

$$A(\chi)(u, v) := A((u, v))\, \chi(s(u, v)) .$$

The signed matrix has $\chi(s(u, v))$ in place of every non-zero entry, i.e., edge (u, v) in X.

We will restrict to analyzing abelian groups G, and for concreteness, the reader can always think of the group as $G = \mathbb{Z}_\ell$, i.e., the cyclic group. Group-based lifts have three properties that make it useful for expander constructions:

1. **Degree** — The degree of the base graph is preserved.
2. **Symmetry** — The group G acts as on the lifted graph as $g\cdot(u, h) = (u, h \cdot g^{-1})$.
3. **Explicit Spectrum** — Using Theorem 1.3.8 and Fact 1.3.9 one can deduce that the eigenvalues of the lifted graph are $\mathrm{Spec}(A(s)) = \bigcup_{\chi \in \hat{G}} \mathrm{Spec}(A(\chi))$. Thus, the lifted graph inherits the spectrum of the base graph (when χ is trivial), and all the new eigenvalues are given by the signed matrices corresponding to non-trivial characters. Therefore, the expansion of the lifted graph is,

$$\lambda(X(s)) \;\le\; \max_{\mathrm{triv} \neq \chi \in \hat{G}} \left\{ \lambda(X),\, \rho(A(\chi)) \right\}.$$

2.1.1 Related Work

Most work in the literature focuses on unstructured lifts (just called lifts) as the goal is to construct expanders without any symmetry requirement. An *unstructured* ℓ-lift is where

each vertex is replaced by ℓ-copies, and for every edge $e = (u, v)$ of X_0, we add a matching between the two sets of ℓ vertices corresponding to u and v.

There are broadly three main techniques to analyze graph lifts, both group-based and unstructured. The techniques, as listed below, give progressively better expansion guarantees but are also increasingly complex and apply in more restricted settings.

Discrepancy Amit and Linial [AL02] introduced random graph lifts in theoretical computer science, and a sequence of works studied various properties of this random model [ALMR01, AL06]. Eventually, Bilu and Linial [BL06] studied the spectral expansion of *2-lifts* giving an explicit construction of graphs with expansion $O(\sqrt{d}\log^{1.5}(d))$ for every degree. Agarwal, Chandrasekaran, Kolla, and Madan [ACKM19] refined the techniques of [BL06], and showed that *random* $\mathbb{Z}_\ell$-lifts (also known as *shift lifts*) are expanding.

Theorem 2.1.3 ([ACKM19, Theorem 1.2])**.** *Let* X_0 *be a* d*-regular* n*-vertex graph, where* $2 \le d \le \sqrt{n/(3\ln n)}$*. Let* X *be a* random $\mathbb{Z}_\ell$*-lift of* X_0*. Then,* $\lambda_u(X) = O(\lambda_u(X_0))$ *with probability* $1 - \ell \cdot e^{-\Omega(n/d^2)}$*. Moreover, for any abelian group* G *such that* $|G| \ge \exp(O_\varepsilon(nd))$*, there does not exist a* G*-lift of* X_0 *such that* $\lambda_u(X) \le \varepsilon \cdot d$*.*

Trace Method The trace method was used by Friedman [Fri03] to prove that random d-regular graphs are near-Ramanujan expanders with high probability. Inspired by a simplified proof of this theorem by Bordenave [Bor20], Mohanty, O'Donnell, and Paredes [MOP20] gave the first explicit construction of near-Ramanujan, i.e., largest non-trivial eigenvalue bounded by $2\sqrt{d-1} + \varepsilon$, graphs of every degree. The key technique in their work was a derandomization of the 2-lifts. Subsequently, Alon [Alo21] gave explicit constructions of near-Ramanujan expanders of every degree and every number of vertices. The work in [MOP20] was also generalized to achieve finer spectral guarantees together with local properties via *unstructured* ℓ-lifts in O'Donnell and Wu [OW20].

Method of Interlacing Polynomial In a breakthrough work, Marcus, Spielman, and Srivastava [MSS15] proved that bipartite graphs admit[1] a $\mathbb{Z}_2$-lift that preserves the spectrum, implying the construction of bipartite Ramanujan expanders. They introduced a new technique called the *method of interlacing polynomials* which has since been extended to analyze G-lifts for $\mathbb{Z}_3, \mathbb{Z}_4, \mathbb{Z}_5$, and various non-abelian groups [HPS16]. One drawback is that this method can only bound λ_2 (or λ_n) but not $\rho_2(A) = \max(\lambda_2, |\lambda_n|)$. Therefore, the applicability of this method (in the context of expanders) is restricted to bipartite graphs.

Overview of the following two chapters We will utilize both methods to prove the existence of an explicit G-lift. One can go through the proof of Theorem 2.1.3 in [ACKM19], and derandomize it directly via expander walks. However, this chapter gives an alternate, simple proof (albeit with a logarithmically weaker bound) via the discrepancy method. This proof easily derandomizes via expander walks when the group is exponentially large in the size of the base graph, i.e., $|G| = 2^{\Omega(n)}$.

In Chapter 3, we turn to the trace method to derandomize smaller lifts. We will extend the techniques of [MOP20]–who analyze 2-lifts–to much larger abelian G-lifts, for $3 \leq |G| \leq 2^{O(n^\delta)}$, i.e., upto subexponential sized lifts.

2.2 Discrepancy Method

The discrepancy method bounds the spectral radius by analyzing the quadratic form $y^\intercal M x$. This quadratic form is approximated by discretizing the space of vectors, i.e., $\mathbb{R}^n$, and considering the form over Boolean vectors i.e., $x, y \in \{0,1\}^n$. The two main steps in this method are: (i) bound the quadratic form over Boolean vectors, and (ii) prove that such a bound implies a bound over the space of real vectors (with some loss). The second step can be seen as a converse to the expander mixing lemma and was proved in a general

1. It is an existential result, it is unknown if random lifts are Ramanujan with high probability.

setting by Bilu and Linial (Theorem 2.2.2). Therefore, this chapter will focus on proving the bound on the discretized quadratic form. We will study two settings, the first with uniformly random signings and the second with signings via *expander walks*. This yields an explicit construction of expanding graphs with exponential lift size.

Theorem 2.2.1 (Exactly Exponential Lifts)**.** *For any positive integer* n *and every constant degree* $d \geq 3$*, given the generating elements of an abelian group* G*, there exists a deterministic* $\mathrm{poly}(\exp(n), |G|)$ *time algorithm that constructs a* d*-regular graph* X *on* $n\ell$ *vertices such that,*

- X *is* G*-lift of a graph* X_0 *on* n *vertices, and*
- *If* $|G| \leq \exp\left(\Theta\left(\frac{n}{\sqrt{d}}\right)\right)$*, then* $\lambda_u(X) \leq O\left(\sqrt{d} \cdot \log d\right)$.
- *If* $|G| = \exp\left(\Theta\left(nd^{\delta}\right)\right)$ *for* $\delta \in [-\frac{1}{2}, 1)$*, then* $\lambda_u(X) \leq O\left(d^{\frac{2+\delta}{3}} \cdot \log d\right)$.

In particular, we have explicit polynomial time construction of a lift when $|G| = \exp(\Theta(n))$*.*

2.2.1 A Simple Proof via Converse EML

In this section, we give a simpler proof of a weaker result similar to one in [ACKM19] that says that if the lifts were picked independently and uniformly at random, the lifted graph also expands. We start with a converse of the expander mixing lemma by Bilu and Linial.

Theorem 2.2.2 ([BL06, Lemma 3.3])**.** *Let* M *be an* $n \times n$ *real symmetric matrix such that the* ℓ_1 *norm of each row in* M *is at most* d*, and all diagonal entries of* M *are, in absolute value,* $O(\alpha(\log(d/\alpha) + 1))$*. Assume that,*

$$\max_{\substack{u,v \in \{0,1\}^n \\ \mathrm{Supp}(u) \cap \mathrm{Supp}(v) = \emptyset}} \frac{|u^t M v|}{\|u\|\|v\|} \leq \alpha.$$

Then, the spectral radius of M *is* $O(\alpha(\log(d/\alpha) + 1))$*.*

We will use $\rho(M)$ to denote the spectral radius of M. Our goal is to bound $\lambda(A(s)) = \max_{\chi \neq 1} A(\chi)$. Since $A(\chi)$ is complex-valued, we will use this simple observation.

Observation 2.2.3 (From complex to real matrices)**.** Let $M = C + iD$ where C, D are real symmetric matrices, then $\rho(M) \leq 2 \cdot \max\{\rho(C), \rho(D)\}$.

Proof. Let $v = v_1 + iv_2$ be an eigenvector and $\alpha = \max\{\rho(C), \rho(D)\}$. Then,

$$\begin{aligned} v^*Av &= \mathrm{Re}(v^*Av) = (v_1^T C v_1 + v_2^T C v_2 - v_1^T D v_2 + v_2^T D v_1) \\ &\leq \rho(C)\|v\|^2 + 2\rho(D)\|v_1\|\|v_2\| \\ &\leq \alpha(\|v_1\| + \|v_2\|)^2 \\ &\leq 2\alpha\left(\|v_1\|^2 + \|v_2\|^2\right) \\ &= 2\alpha\|v\|^2. \end{aligned}$$

□

Let X be a graph, G be an abelian group, and s be a signing $s : E \to G$. Let $A(\chi) = C + iD$ where C, D are real symmetric matrices. Now, we need to bound the discrepancy of C, D.

The argument is as follows: Fix a pair of vectors u, v, and a character χ. We show that if the signing is (pseudo)random, then the discrepancy (with respect to u, v) is small with a very high probability. We then take a union bound over the 2^{2n} pairs of boolean vectors and $|G| - 1$ many non-trivial characters. To carry this out, we need an exponentially good probability; hence, we will need a Hoeffding bound. We start with a simple calculation.

Some useful inequalities Let S, T be subsets of the vertices of a d-regular graph. Define $E(S, T) = \{(x, y) \in E \mid x \in S,\ y \in T\}$ and let $e(S, T) := |E(S, T)|$. Let $u, v \in \{0, 1\}^n$ and let $S := \mathrm{Supp}(u), T := \mathrm{Supp}(v)$. Then,

$$|u^T C v| \leq \sum_{u \in E(S,T)} |\mathrm{Re}(\chi(s(e)))| \leq e(S, T), \tag{2.1}$$

$$|u^T D v| \leq \sum_{u \in E(S,T)} |\mathrm{Im}(\chi(s(e)))| \leq e(S, T). \tag{2.2}$$

Let us now state the expander mixing lemma,

$$\left| e(S, T) - \frac{d|S||T|}{n} \right| \leq \lambda_u(X)\sqrt{|S||T|}. \tag{2.3}$$

We state the key property we require from our distribution over the signings. For a complex number z, we use $\mathrm{Re}(z), \mathrm{Im}(z)$ to refer to its real and imaginary parts respectively.

Definition 2.2.4 (β-exponentially good signings)**.** Fix a graph X and a group G. A distribution $\mathcal{D}$ over signings $s : E \to G$ is β-exponentially good if for any subset of edges, $U \subseteq E$, and any non-trivial character $\chi : G \to \mathbb{C}$,

$$\Pr_{s\sim\mathcal{D}}\left[\left|\sum_{e\in U} \mathrm{Re}(\chi(s(e)))\right| \geq t\right] \;\leq\; 2\exp\left(\frac{-\beta t^2}{|U|}\right),$$

$$\Pr_{s\sim\mathcal{D}}\left[\left|\sum_{e\in U} \mathrm{Im}(\chi(s(e)))\right| \geq t\right] \;\leq\; 2\exp\left(\frac{-\beta t^2}{|U|}\right).$$

This property holds for uniformly random signings and pseudorandom signings using walks on an expander graph (Corollary 2.2.10). We only prove the general case later.

Lemma 2.2.5 (Hoeffding)**.** *Let $\mathcal{D}$ be the uniform distribution over all signings $s : E \to G$. Then $\mathcal{D}$ is $\frac{1}{128e}$-exponentially good.*

Lemma 2.2.6. *Let s be sampled from a β-exponentially good distribution, and let N be either C or D, where these are the matrices defined above. Let $\gamma^3 = \frac{\sqrt{d}}{2\beta n}\ln(3|G|)$, and define $\alpha = (\gamma+1)\lambda$. Then for every pair of vectors $u, v \in \{0,1\}^n$, $|u^\intercal N v| \leq \alpha\|u\|\|v\|$, except with probability $\frac{2}{3|G|}$.*

Proof. Since the proofs are identical, we use N as a placeholder, which can be replaced by C or D. Let $S := \mathrm{Supp}(u), T := \mathrm{Supp}(v)$ and define $a := \sqrt{\|u\|\|v\|} = \sqrt{|S||T|}$.

Case 1 - $a \leq \frac{\gamma n\lambda}{d}$. From Eq. (2.1) and Eq. (2.3), we have,

$$|u^\intercal N v| \;\leq\; e(S,T) \;\leq\; \frac{d}{n}a^2 + \lambda a \;\leq\; (\gamma+1)\,a\lambda.$$

Case 2 - $a > \frac{\gamma n\lambda}{d}$. Using the trivial bound that $a \leq n$ in Eq. (2.3), we get,

$$e(S,T) \;\leq\; a\left(\frac{da}{n} + \lambda\right) \;\leq\; a(d+\lambda) \leq 2ad.$$

By Lemma 2.2.5 we get,

$$\Pr_{s\sim\mathcal{D}}\left[|u^\intercal N v| \geq (\gamma+1)\,a\lambda\right] \;\leq\; 2\exp\left(\frac{-((\gamma+1)a\lambda)^2}{128e\cdot e(S,T)}\right).$$

We can upper bound this as,

$$\begin{aligned}
\frac{((\gamma+1)a\lambda)^2}{128e\cdot e(S,T)} &\geq \frac{((\gamma+1)a\lambda)^2}{128e(2ad)} \\
&\geq \frac{(\gamma+1)^2 a}{256e} && (\lambda^2 > d) \\
&\geq \frac{\gamma(\gamma+1)^2 n\lambda}{256ed} && \text{(By case assumption on } a\text{)} \\
&\geq \frac{\gamma^3 n}{2\cdot 256e\sqrt{d}} \\
&\geq \ln(3|G|). && \text{(By assumption on } \gamma\text{),}
\end{aligned}$$

□

Theorem 2.2.7 (Random signings)**.** *Let X be a d-regular graph that is a $\lambda = O(\sqrt{d})$-expander. Let $\mathcal{D}$ be a β-good distribution. Then, there exists a signing $s \in \mathrm{Supp}(\mathcal{D})$ such that the lifted graph, $X(s)$, is a λ'-expander, where $\lambda' = O\big(\max\big\{d^{\frac{2+\delta}{3}}, \sqrt{d}\big\}\log d\big)$.*

Proof. Lemma 2.2.6 gives a bound of $\alpha = (\gamma+1)\lambda$ on the Rayleigh quotient of C, D holds except with probability $\frac{2}{3|G|}$ over the signings. Since $\lambda = O\big(\sqrt{d}\big)$, we get $\alpha = O\big(\max\big\{d^{\frac{2+\delta}{3}}, \sqrt{d}\big\}\big)$.

Since the graph is d-regular, $A(\chi)$ is d-sparse and so is C and D. The ℓ_1-norm of any row of $C, D \leq d$ as we have a sum of d entries of the form $\mathrm{Re}(\omega^j), \mathrm{Im}(\omega^j)$ for some j and the absolute value of each of these is upper bounded by 1. Moreover, the diagonal entries are all zero. Therefore, C, D satisfy the criteria of the Theorem 2.2.2 which implies,

$$\lambda_{\max} A(\chi) \leq 2\max\{\rho(C), \rho(D)\} \leq 2\alpha\log(d/\alpha) \leq O(\alpha\log d).$$

To finish the proof, we need to take a union bound over each of the $\ell - 1$ non-trivial characters and bound the spectrum of $A(\chi)$ as above. Thus, we have that the probability that there exists a good signing is at least $1 - \ell\left(\frac{2}{3\ell}\right) > 0$. □

2.2.2 *Derandomized signings via expander walks*

Now, we describe how the signings for the lift are generated via walks on an expander. The key pseudorandom property is that these signings satisfy an expander Hoeffding bound, i.e., are β-exponentially good. Let $\ell = |G|$, and assume that we have a numbering of the group elements, $\varphi : [\ell] \to G$. We use an expander construction from Alon [Alo21],

Theorem 2.2.8 ([Alo21, Thm. 1.3])**.** *For every degree $d \geq 3$, every $\varepsilon > 0$ and all sufficiently large $m \geq n_0(d, \varepsilon)$ where md is even, there is an explicit construction of a (d, λ)-expander graph on m vertices with $\lambda_u \leq 2\sqrt{d-1} + \varepsilon$.*

We can fix the degree to be an even constant, say d', and have $\varepsilon = \sqrt{d'-1}$. Then, we use Theorem 2.2.8 to get an explicit expander, L, on ℓ vertices. To obtain a sequence of lifts, i.e., elements of G, we first pick a random vertex, v_1, which uses $\log \ell$ bits of randomness. Then, we do a random walk for $dn - 1$ steps producing a sequence $(v_1, \cdots, v_{dn})$ of vertices of L, which we interpret as elements of G as $(\varphi(v_1), \cdots, \varphi(v_{dn}))$. Each random walk step requires $O(\log d')$ bits of randomness as the graph is d'-regular. Therefore, the total amount of randomness is[2] $O(\log \ell + (dn-1)\log d')$. The main observation is that signings generated via expander random walks satisfy a Hoeffding-type concentration result we used in the earlier proof.

Theorem 2.2.9 ([Rao19, Thm. 1.1])**.** *Let $\{Y_i\}_{i=1}^{\infty}$ be a stationary Markov chain with state space $[N]$, transition matrix A, stationary probability measure π, and averaging operator E_π, so that Y_1 is distributed according to π. Let $\lambda = \|A - E_\pi\|_{L_2(\pi)\to L_2(\pi)}$ and let $f_1, \cdots, f_t : [N] \to \mathbb{R}$ so that $\mathbb{E}[f_i(Y_i)] = 0$ for all i and $|f_i(v)| \leq a_i$, for all $v \in [N]$ and all i. Then for $u \geq 0$,*

$$\Pr\left[\left|\sum_{i=1}^{t} f_i(Y_i)\right| \geq u\left(\sum_{i=1}^{t} a_i^2\right)^{\frac{1}{2}}\right] \;\leq\; 2\exp\left(\frac{-u^2(1-\lambda)}{64e}\right).$$

Using the above concentration result, we can prove a general version of Lemma 2.2.5.

2. Another way to say this is that the number of walks of length on dn on L is $|G| \cdot d'^{dn}$.

Corollary 2.2.10 (Expander Hoeffding)**.** *Let $\mathcal{D}$ be the uniform distribution over signings constructed by a random walk on the above-described expander. Then, $\mathcal{D}$ is $\frac{1}{128e}$-exponentially good.*

Proof. Let $Y_e = s(e)$ be the random variables associated with each edge e. From the construction described earlier, $\{Y_{u,v}\}$ is a Markov chain with the transition matrix being the weighted adjacency matrix of the expander L such that $\lambda(L) \leq 3\sqrt{d'-1}/d'$. Thus, $1-\lambda \geq 1 - \frac{3}{\sqrt{d'}} \geq 1/2$ for $d' \geq 36$. The stationary measure π is the uniform measure on vertices of L, and it is stationary as the all-ones vector is an eigenvector of the weighted adjacency matrix with eigenvalue 1. Recall that we picked the first vertex (Y_1) uniformly, i.e., from π. Let $f_e = \mathrm{Re}(\chi(s(e)))$ and $g_e = \mathrm{Im}(\chi(s(e))$ if $e \in U$ and 0 otherwise.

$\mathbb{E}[f_e] = \frac{1}{\ell}\sum_{i=0}^{\ell-1} \mathrm{Re}(\omega^i)$ because the characters are roots of unity and the expectation is over π which is uniform. Since the sum of the roots of unity is zero, so is its real and imaginary part. This holds thus for g_e too. Moreover, $a_e = 1$ if $e \in U$ and is 0 otherwise. Applying Theorem 2.2.9 with $u := t/\sqrt{|U|}$ gives the result. □

Theorem 2.2.1 (Exactly Exponential Lifts)**.** *For any positive integer n and every constant degree $d \geq 3$, given the generating elements of an abelian group G, there exists a deterministic $\mathrm{poly}(\exp(n), |G|)$ time algorithm that constructs a d-regular graph X on $n\ell$ vertices such that,*

- *X is G-lift of a graph X_0 on n vertices, and*
- *If $|G| \leq \exp\left(\Theta\left(\frac{n}{\sqrt{d}}\right)\right)$, then $\lambda_u(X) \leq O\left(\sqrt{d}\cdot\log d\right)$.*
- *If $|G| = \exp\left(\Theta\left(nd^{\delta}\right)\right)$ for $\delta \in [-\frac{1}{2}, 1)$, then $\lambda_u(X) \leq O\left(d^{\frac{2+\delta}{3}}\cdot\log d\right)$.*

In particular, we have explicit polynomial time construction of a lift when $|G| = \exp(\Theta(n))$.

Proof. We construct a d-regular graph X_0 using Theorem 2.2.8 on n vertices such that $\lambda_2(G) \leq 2\sqrt{d}$. We generate a set of signings as described above using a d'-regular expander on ℓ vertices. This takes time $\ell\exp(nd\ln(d'))$ and we can fix $d' = 36$. For each signing, we compute the eigenvalue of the adjacency matrix of the lifted graph and pick the one with the smallest λ_2. The existence of a good signing is guaranteed by Theorem 2.2.7. □

CHAPTER 3

GRAPH LIFTS VIA THE TRACE METHOD

Janki: Derandomizing union bounds is hard!

Piro: Yeah, they are tricky. It is often helpful to write an alternate proof.

Laqa: The random lift analysis via discrepancy I mentioned last time has the same issue. I am reading this trace method proof of random 2-lifts that avoids a union bound.

Janki: What does the trace method do?

Laqa: Think of it as a reduction to counting! Basically, $\rho(A)^k \leq \mathrm{tr}(A^k)$ but trace counts cycles.

Piro: So now, instead of using edge expansion of the base graph, I guess you want to use the fact that it has a few short cycles. Where is the randomness used?

Laqa: Yes, precisely. The randomness zeroes out many cycles, and for the expected trace, you only need to count certain special cycles that are easier to bound.

Janki: I need details. Care to explain while we wait an hour for our deep dish?

Laqa: Sure! So let A be the adjacency matrix of an expander ...

In this chapter, we continue our analysis of random lifts. We use the trace power method and prove the following result for smaller lifts,

Theorem 3.0.1. *For large enough* n *and constant degree* $d \geq 3$*, given an Abelian group* G *such that* $|G| \leq \exp(n^{\Theta(1)})$*, and any fixed constant* $\varepsilon \in (0,1)$*, we can construct in deterministic polynomial time, a* d*-regular graph* X *on* $\Theta(n\ell)$ *vertices such that,*

- X *is* G*-lift of a graph* X_0 *on* $\Theta(n)$ *vertices.*
- *If* $|G| \leq \exp\left(n^{\delta(d,\varepsilon)}\right)$*, then* $\lambda(X) \leq 2\sqrt{d-1} + \varepsilon$*.*
- *If* $|G| \leq \exp\left(n^{\delta}\right)$ *and also* $d \geq d_0(\varepsilon)$*, then* $\lambda(X) \leq \varepsilon \cdot d$*.*

3.1 Overview of the Trace Method

The trace method is the name for utilizing the following inequality,

$$\rho(B)^{2k} \;\leq\; \sum_i |\lambda_i^k|^2 \;=\; \mathrm{tr}((B^*)^k B^k) \quad .$$

The LHS is the quantity we wish to bound, while the RHS can be reduced to a combinatorial quantity for certain operators B. We will switch to working with the non-backtracking operator $B_s(\chi)$ instead of the adjacency matrix $A(\chi)$ as the combinatorial quantity arising from the trace of B^k is easier to work with.

Definition 3.1.1 (Non-backtracking walk operator)**.** For an extended signing $s : E^d \to G$ and a character χ of G, the signed non-backtracking walk matrix $B_s(\chi)$ is a non-symmetric matrix of size $|E^d| \times |E^d|$ in which the entry corresponding to the pair of edges $(u,v),(x,y)$ is $\chi(s(x,y))$ if $v = x,\ u \neq y$, and zero otherwise.

The unsigned variant is obtained by taking the trivial character in the definition above. Let the non-backtracking walk matrix of X be B and the lifted graph with respect to a signing s be $B_{\chi(s)}$. We use the following standard facts that we prove in Appendix A.2.

Fact 3.1.2. *Let* B *be the non-backtracking walk matrix of a* d*-regular graph* G*. Then,*

$$\lambda(G) \;\leq\; 2 \cdot \max\left\{\sqrt{d-1},\ \rho_2(B)\right\}.$$

Fact 3.1.3. *If* G *is abelian, then* $\mathrm{Spec}(B_{\chi(s)}) = \bigcup_{\chi \in \hat{G}} \mathrm{Spec}(B_s(\chi))$.

The overall approach is as follows: (i) Show that over random signings, the quantity $\mathbb{E}_s[\mathrm{tr}((B_s^*)^k B_s^k)]$ is small (for every non-trivial character) (ii) Derandomize part (i) by ob-

serving that one only requires k-wise independence rather than uniformly independent signings. This yields an explicit construction of a signing. We now explain part (i) in detail as it is the technical crux of the proof.

3.2 Proof strategy

Let X_0 be a base expander graph and $s\colon E_0 \to \mathbb{Z}_2$ be a signing that defines a lift. It is convenient to first think that the signing is chosen uniformly at random, and later, see which properties were used so that an appropriate derandomization tool may be used.

The [MOP20] Argument: Applying the trace method[1] we get,

$$\rho(B_s)^{2k} \;\le\; \mathrm{tr}((B_s^*)^k B_s^k) \;=\; \sum_{\substack{(e_1,\ldots,e_{2k})\\ \text{closed edge walk}}} \prod_{i=1}^{2k} \chi(s(e_i)).$$

The above expression greatly simplifies when we take the expectation over a uniformly random signing since only walks in which every edge occurs at least twice do not zero out on average. These walks are called *singleton-free* in [MOP20]. We have,

$$\mathop{\mathbb{E}}_{s}\left[\rho(B_s)^{2k}\right] \;\le\; \sum_{\substack{(e_1,\ldots,e_{2k})\\ \text{closed edge walk}}} \mathop{\mathbb{E}}_{s}\left[\prod_{i=1}^{2k} \chi(s(e_i))\right] \;\le\; \left|\left\{\begin{array}{c}\text{2k-length singleton-free}\\ \text{non-backtracking walks in } X_0\end{array}\right\}\right|.$$

This reduces the problem of bounding the spectral radius to a counting problem of these special walks. In the hypothetical (idealized) scenario of X_0 being Ramanujan and the counting on the RHS above being $(d-1)^k$, we would have a Ramanujan lift.

One of the main technical contributions in [MOP20] is the counting of $2k$-length singleton-free non-backtracking walks in X_0, which they call *hikes*. For the sake of in-

1. To avoid discussing unimportant technicalities, we will make some simplifications in this high-level overview.

tuition, we will assume that X_0 has girth g, but it is not hard to modify the argument when X_0 has at most one cycle around any neighborhood of radius $< g/2$ centered at a vertex in X_0 (this is the property of being *bicycle-free*). They view the vertices and edges visited in a hike as forming a hike graph $\mathcal{H}$. Assuming that $g = \Omega(\log_{d-1}(n))$, if k is not too large, then $\mathcal{H}$ looks like a tree, possibly with a few additional edges forming cycles as established by Alon, Hoory, and Linial in [AHL02] (and generalized in [MOP20] to bicycle-free radius from girth).

Assuming that the hike is singleton-free, we can have at most k steps that visit an edge that was not previously visited. This implies that the hike graph $\mathcal{H}$ has at most k edges and at most $k+1$ vertices (since it is connected). They count the number of these special walks by directly specifying an encoding for the hike. Up to negligible factors (after $2k$-th root for k not too small), they show that there are at most

$$n \cdot (d-1)^k \cdot k^{O\left(\frac{\ln(k)}{g}\right) \cdot k},$$

singleton-free hikes of length $2k$ (see [MOP20, Theorem 3.9] for precise details). This bound trivializes, i.e., it becomes at least $(d-1)^{2k}$, for $\ln(k) \gg \sqrt{g} = \Theta\left(\sqrt{\log_{d-1}(n)}\right)$. This means that we cannot use their bound for very long walks, and this, in turn, prevents us from getting lift sizes larger than $2^{2^{\Theta(\sqrt{\log_{d-1}(n)})}}$ from their results.

Our Approach Now, let us consider abelian G-lifts and let $\ell := |G|$. The spectral radius of each $B_s(\chi)$ can be analyzed similarly via the trace power method. However, we need to bound all of them *simultaneously*. We know no better way than a simple union bound over the $\ell - 1$ cases, but this will force us to obtain a much better concentration guarantee out of the trace power method, which entails considering much larger walk lengths.

Instead of encoding a hike directly as in [MOP20], we will first encode the subgraph of X_0 traversed by the hike, which we call the hike graph, and then encode the hike having the full hike graph at our disposal. We will give two different encodings for the hike

graph. The first one is simpler and can encode an arbitrary graph. The second encoding uses the special structure of the hike graph, namely, having few vertices of degree greater than 2. Both encodings are based on the traversal history of the simple depth-first search (DFS) algorithm. Let $\mathcal{H}$ be the hike graph on $m \leq k$ edges and $n' \leq k+1$ vertices. As DFS traverses $\mathcal{H}$, each of its edges will be visited twice: first "forward" via a recursive call and later "backwards" via a backtracking operation. We view each step of the DFS traversal as being associated with an edge that is being currently traversed and the associated type of traversal: recursive (R) or backtracking (B). A key observation is that only for the recursive traversals we need to know the next neighbor out of $d-1$ possibilities (except for the first step). For the backtracking steps, we can rely on the current stack of DFS. Thus, if we are given a starting vertex from X_0, a binary string in $\{R, B\}^{2m}$ and a next neighbor for each recursive step, we can reconstruct $\mathcal{H}$. Note that there at most

$$n \cdot d \cdot (d-1)^k \cdot 2^{2k},$$

such encodings. Having access to the hike graph and again assuming that the graph has girth $g = \Omega(\log_{d-1}(n))$ (similarly, bicycle freeness is also enough). Using the locally tree-like structure, a $2k$-length hike can be specified by splitting it into segments of length $< g/2$; by specifying the starting vertex of the first segment and the ending vertex of each segment, we have enough information to recover the full hike. Note that there are at most

$$k^{O(k/g)}$$

ways of encoding a hike. Then, the number of $2k$-hikes in X_0 is at most

$$n \cdot d \cdot (d-1)^k \cdot 2^{2k} \cdot k^{O(k/g)}.$$

Now we can take $k \approx n^{\delta}$ for a sufficiently small $\delta = \delta(d) > 0$ and obtain, after taking the $2k$-th root of the above quantity,

$$\rho(B_s) \;\leq\; (1+\varepsilon) \cdot 2 \cdot \sqrt{(d-1)},$$

when $k = k(n, d, \varepsilon)$ is sufficiently large and $c = c(\varepsilon)$ is sufficiently small. The extra factor

of 2 prevents us from obtaining near-Ramanujan bounds with this counting. Nonetheless, the simple counting already allows us to obtain expansion $O(\sqrt{d})$ for lifts sizes as large as $2^{n^{\delta(d)}}$. Moreover, by weakening the expansion guarantee we can obtain lift sizes as large as $2^{n^{\Theta(1)}}$ from this counting and obtain part of Theorem 3.0.1. If we insist on getting a near-Ramanujan bound, we need to compress the traversal history further since storing a string $\{R, B\}^{2m}$ is too costly and leads to this factor of 2. Note that this string has an equal number of R and B symbols, so it cannot be naively compressed.

To obtain a near-Ramanujan graph, we will take advantage of the special structure of the hike graph (when the walk length is large but not too large) in which most of its vertices have degree 2. These degree 2 vertices are particularly simple to handle in a DFS traversal. For them, we only need to store the next neighbor out of $d - 1$ possibilities in X_0 (except for the first step). In a sequence of backtrackings if the top of the DFS stack is a degree 2 vertex, we know that we are done processing it since no further recursive call will be initiated from it. Then, we simply pop it from the stack. It is for the "rare" at most $\delta \cdot n'$ vertices v of degree ≥ 3 that we need to store how many extra recursive calls t_v we issue from v and a tuple of additional next neighbors $(d_1, \ldots, d_{t_v})$. The total number of such encodings is at most

$$n \cdot d \cdot (d-1)^k \cdot \binom{k+1}{\delta(k+1)} \cdot (d-1)^{\delta(k+1)},$$

which combined with the same previous way of encoding a hike given its graph results in a total number of hike encodings of X_0 of at most

$$n \cdot d \cdot (d-1)^k \cdot \binom{k+1}{\delta(k+1)} \cdot (d-1)^{\delta(k+1)} \cdot k^{O(k/g)},$$

By choosing $\delta = \delta(d, \varepsilon)$ sufficiently small and taking $k = k(n, d, \varepsilon) \leq 2^{\delta \cdot g} \approx n^{O_d(\delta)}$ sufficiently large, we obtain after taking the $2k$-th root

$$\rho(B_s) \leq \sqrt{(d-1)} + \varepsilon,$$

leading to a near-Ramanujan bound for lifts as large as 2^{n^δ} in Theorem 3.0.1.

Now we briefly explain how to handle the union bound to ensure that $\rho(B_s(\chi))$ is *simultaneously* small for all $(\ell-1)$ non-trivial characters (in the decomposition of Fact 3.1.3). This union bound is *standard* when using the trace power method, what is relevant is the trade-off between lift size and walk length. To obtain a high probability guarantee from a guarantee on expectation; it is standard to consider larger walk lengths from which concentration follows from a simple Markov inequality. More precisely, if for some function f, $\mathbb{E}_s[\rho(B_s(\chi))^{2k}] \leq f(n, d, g, k)$, then by Markov's inequality,

$$\Pr_s\left[\rho(B_s(\chi)) \geq 2^{\log_2(\ell)/(2k)} \cdot f(n, d, g, k)^{1/(2k)}\right] \;\leq\; \frac{1}{\ell}.$$

Therefore, for $k \geq \log_2(\ell)$ sufficiently large, we can union bound over all characters χ and obtain similar bounds as before. As alluded to above, this lower bound on the length of the walk depending on the lift size is the reason why we are led to consider much longer walks. To conclude this proof sketch, we need to replace a random signing by a pseudorandom random one. As in [MOP20], we use ε-biased distributions, e.g., the one[2] by Jalan and Moshkovitz in [JM21]. We may be taking very large walks on the base graph X_0, so the error of the generator needs to be smaller than $n \cdot d^{2k}$, where k can be as large as $n^{\Theta(1)}$. We note that as long as the degree d is a constant, this quantity is, at most, a polynomial in the size of the *final* lifted graph X since walks of length $O(\log(|V(X)|))$ suffice for any lift size up to full extent of $2^{O(n)}$, for which abelian lifts can be expanding.

3.3 A New Encoding for Special Walks

In this section we will count the total number of singleton-free hikes of a given length on a fixed graph, X. We split the count into two parts. First, we count the number of possible

2. For our application, it suffices to have the support size of the ε-biased distribution polynomial in $1/\varepsilon$.

hike graphs and then, for a given hike graph $\mathcal{H}$, we count the number of hikes that can i.e., yield $\mathcal{H}$ on traversal. Each of these counts is via an encoding argument and therefore we have two kinds of encoding. One for graphs and the other for hikes. In the first part of the section we give two ways of encoding graphs, and in the other half, we encode hikes. Since the first section is a general encoding for subgraphs, we relegate formal definitions related to hikes to a later section.

3.3.1 Graph Encoding

Let $\mathcal{H}$ be a subgraph of a fixed d-regular graph X. We wish to encode $\mathcal{H}$ in a succinct way such that given the encoding and X, we can recover $\mathcal{H}$ uniquely. We will give two ways of encoding $\mathcal{H}$. The first one will be generic that works for any subgraph of a d-regular graph. The second encoding takes advantage of the special sparse structure (not too many vertices of degree greater than two). We assume that we have an order on the neighbors of every vertex, and thus, given (v, j), we can access the j^{th} neighbor of v efficiently.

We will do this by encoding a DFS based-traversal of it from a given start vertex . Here, we really need our DFS traversal to be optimal in the sense that the number of times each edge is traversed is at most two and not any higher. We, therefore, include precise details of our implementation in Appendix A.3. To reconstruct the graph, we reconstruct the traversal and so need access to two types of data before every step:

1. Is this step recursive or backtracking?

2. If it is a recursive step, then which neighbor do we recurse to?

To determine the neighbor of the current vertex, we need to move to in a recursive call, we need to specify one out of the $d - 1$ possibilities (except in the first step, which has d possibilities). This can be specified by a tuple of $(d_1, \ldots, d_{|E(\mathcal{H})|}) \in [d] \times [d-1]^{|E(\mathcal{H})|-1}$ indicating the neighbor. For a backtracking step, we just pop the stack and thus do not

need any additional data.

We use two ways to figure out whether a step is recursive or backtracking. The direct way is to just record the sequence in a binary string of length $2|E(\mathcal{H})|$. A neighbor u of v is called *recursive* if the edge (v, u) is visited by a recursive call from v. A simple observation about backtracking sequences is that – a backtracking sequence starts when we encounter a vertex that has already been visited or when we reach a degree one vertex. The sequence ends when we see a visited vertex that has unvisited recursive neighbors. Therefore, we store a string $\sigma \in [d] \times [d-1]^{|V(\mathcal{H})|-1}$ in which σ_i denotes the number of recursive neighbors of the i^{th} visited vertex. To summarize,

GraphEnc($\mathcal{H}$):

(a) Starting vertex $v_1 \in V(G)$

(b) A sequence of degrees $(d_1, \ldots, d_{|E(\mathcal{H})|}) \in [d] \times [d-1]^{|E(\mathcal{H})|-1}$

(c) Either $\sigma \in \{R, B\}^{2|E(\mathcal{H})|}$ (**Encoding I**) or,
$\sigma \in [d] \times [d-1]^{|V(\mathcal{H})|-1}$ (**Encoding II**)

Algorithm 3.3.1 (StepType)**.**

Input (v, t)

Output (Type)

Note - The subroutine to detect the type of step depends on the encoding string σ.

· If σ is from **Encoding I**, return σ_t

· Else, let $j = \mathrm{ord}(v)$

  · If $\sigma_j > 0$ //Check if there are any remaining recursive neighbours

    · Decrement $\sigma_j \leftarrow \sigma_j - 1$

    · return R

  · Else, return B

Algorithm 3.3.2 (Unpacking Algorithm for **GraphEnc**).

Input **GraphEnc**($\mathcal{H}$)

Output $\mathcal{H}$

- Initialize DFS stack S with v_1
- Initialize $\mathcal{H} = (\{v_1\}, \emptyset)$
- Initialize $n, r, t = 1$ // count visited vertices, recursive steps and total steps
- Initialize $\mathrm{ord}(v_1) = 1$
- While $S \neq \emptyset$:
    - Let v be the top vertex on the stack S
    - step = StepType(v, t)
    - If step = R (recursive):
        - Assign v_{next} to be d_r^{th} neighbor of v and increment r
        - Add edge $\{v, v_{\text{next}}\}$ to $\mathcal{H}$
        - If v_{next} is unvisited :
            - Add vertex v_{next} to $\mathcal{H}$
            - $n \leftarrow n + 1$
            - $\mathrm{ord}(v_{\text{next}}) \leftarrow n$
            - push(v_{next}, S)
        - Else if v_{next} is visited, increment t // Next step is backtracking
    - If step = B (backtracking):
        - pop(S)
        - $t \leftarrow t + 1$
- return $\mathcal{H}$

Counting the encodings

For the first kind of encoding of type, we have 2^{2k} strings of length $2k$ over $\{R, B\}$. The second encoding might seem wasteful in general but it is much better when the graph has special structure that our hike graph will satisfy. We first note that for any vertex v, the number of recursive neighbours $\sigma_v \leq \deg_{\mathcal{H}}(v) - 1$ (or $\leq \deg_{\mathcal{H}}(v)$ if $v = v_0$).

Definition 3.3.3 (Excess)**.** The *excess* of $\mathcal{H}$ is defined as $\mathrm{exc}(\mathcal{H}) := |E(\mathcal{H})| - |V(\mathcal{H})|$.

Definition 3.3.4 (Excess Set)**.** We define a vertex to be an *excess vertex* in $\mathcal{H}$ if $\deg_{\mathcal{H}}(v) > 2$ and we define the *excess set* to be the set consisting of such vertices, i.e.,

$$\mathrm{excSet}(\mathcal{H}) := \left|\left\{v \in V(\mathcal{H}) \mid \deg(v) > 2\right\}\right|.$$

Lemma 3.3.5. *Let X be a fixed d-regular graph on n vertices. The total number of connected subgraphs $\mathcal{H}$ of X having at most $\leq k$ edges is at most*

$$2n \cdot d \cdot (d-1)^{k-1} \cdot 2^{2k}.$$

Moreover, if $\mathcal{H}$ is constrained to have at most two vertices of degree one[3] *and $\mathrm{exc}(\mathcal{H}) \leq \delta k$, the count is at most*

$$2nk^3 \cdot d \cdot (d-1)^{k-1} \cdot 2^{H_2\left(\frac{\delta}{1-\delta}\right)k} \cdot d^{\delta k}.$$

Proof. We first fix the number of edges as m and we will then sum up the expression for $m \leq k$. Algorithm 3.3.2 unambiguously recovers the graph and therefore the number of possible graphs can be counted by counting the number of possible inputs. The number of degree sequences and start vertices are $n \cdot d(d-1)^{m-1}$. The number of σ-strings of encoding I are 2^{2m}. Therefore for a given m, we have $nd \cdot (d-1)^{m-1} \cdot 2^{2m}$ and summing this gives the first claim.

In the second case, the key idea is that for every vertex (except the start) of degree 2, σ_v must be 1. Since $|\mathrm{excSet}(\mathcal{H})| \leq \delta m$, almost all of the string σ is filled by 1.

3. We will see later that hike graphs satisfy this strange property

We first pick the number of vertices, say t. There are at most m choices for this. Then, we let the number of excess vertices be j. Summing over all possible j, the number of σ-strings of length t is at most

$$t^2 \sum_{j=0}^{\delta m} \binom{t}{j} d^j \;\leq\; t^2 d^{\delta m} \sum_{j=0}^{\delta m} \binom{t}{j} \;\leq\; t^2 d^{\delta m} 2^{H_2\left(\frac{\delta}{1-\delta}\right)t}.$$

Here, the first term counts the ways of having or up to two vertices of degree 1, the second counts the ways to choose the excess vertices, and the third counts the number of their recursive neighbors. In the last inequality, we used that $t = m - exc(\mathcal{H}) \geq (1-\delta)m$.

The complete expression for the number of graphs would then be

$$\sum_{m \leq k} \left(nd(d-1)^{m-1} \sum_{t=(1-\delta)m}^{m} t^2 d^{\delta m} 2^{H_2\left(\frac{\delta}{1-\delta}\right)t} \right) \;\leq\; 2nk^3 \cdot d \cdot (d-1)^{k-1} \cdot 2^{H_2\left(\frac{\delta}{1-\delta}\right)k} \cdot d^{\delta k}.$$

□

3.3.2 Bounding Singleton Free Hikes

Following [MOP20], we make the following useful definitions,

Definition 3.3.6 (Singleton-free hikes)**.** A k-*hike* W is a closed walk of $2k$-steps[4] in X in which every step except possibly the $(k+1)^{st}$ is non-backtracking. A hike is *singleton-free* if no edge is traversed exactly once.

Definition 3.3.7 (Bicycle free radius [MOP20])**.** A graph X is said to have a bicycle-free radius at radius r if the subgraph $\mathcal{H}$ of distance-r neighborhood of every vertex has $exc(\mathcal{H}) \leq 0$.

We will work with singleton-free hikes in this section. A singleton-free k-hike on X defines a subgraph $\mathcal{H}$ such that there at most two vertices of degree 1 (the start vertex and

4. That is sequence of $(v_0, \cdots, v_{2k-1})$ such that $(v_i, v_{i+1}) \in E(G)$ and $v_0 = v_{2k-1}$

the middle vertex) and the number of edges is at most k as every edge is traversed at least twice. The goal now is to count the possible number of singleton-free k-hikes that yield a fixed subgraph $\mathcal{H}$. Having access to $\mathcal{H}$, we will need to encode the hike in a way similar to the encoding of stale stretches in [MOP20].

HikeEnc:

(a) $(v_1, \ldots, v_s) \in V(\mathcal{H})^s$,where $s = \lceil 2k/r \rceil$ and r is the bicycle free radius of $\mathcal{H}$.

(b) $(c_1, \ldots, c_s) \in \{0, \pm 1, \cdots, \pm\lfloor r/2 \rfloor\}^s$. Here, c_i denotes the number of times the unique cycle (in the neighborhood of v_i) is to be traversed and the sign indicates the orientation. Since each stretch is of length r and each cycle of length at least 2 we can traverse a cycle at most $\lfloor r/2 \rfloor$ times.

Claim 3.3.8. *For any graph $\mathcal{H}$ that is bicycle free at radius r, the number of simple singleton-free k-hikes that have $\mathcal{H}$ as their hike graph is at most* $(|rV(\mathcal{H})|)^{\lceil 2k/r \rceil}$.

Proof. Follows from the possible values the encoding **HikeEnc** can take. □

We use a generalization of the bound of Alon et al. [AHL02] on the excess number (originally involving the girth), extended to bicycle-free radius in [MOP20].

Theorem 3.3.9 ([MOP20, Theorem 2.13])**.** *Let $\mathcal{H}$ be a bicycle free graph of radius $r \geq 10 \ln(|V(\mathcal{H})|)$. Then,*

$$\mathrm{exc}(\mathcal{H}) \leq \frac{\ln(e|V(\mathcal{H})|)}{r} \cdot |V(\mathcal{H})|.$$

Corollary 3.3.10. *Let X be a d regular graph on n vertices bicycle free at radius r. Let $\mathcal{H}$ be a subgraph with at most two vertices of degree one on n_0 vertices where $n_0 = e^{\delta r - 1}$ for some $\delta \leq 1/10$. Then,*

$$\mathrm{excSet}(\mathcal{H}) \leq 2\delta n_0 + 2.$$

Lemma 3.3.11. *Let X be a d regular graph, with $d \geq 3$, on n vertices bicycle free at radius r. Then, the total number of singleton-free $(k-1)$-hikes on X is at most*

$$\left(2^{\gamma_1}\sqrt{d-1}\right)^{2k} \text{ where } \gamma_1 = 1 + \frac{\log(nrk)}{2k} + \frac{\log(rk)}{r}.$$

If we assume that $3 \le k \le e^{\delta r}$*, then it is at most*

$$\left(2^{\gamma_2}\sqrt{d-1}\right)^{2k} \text{ where } \gamma_2 = \frac{\log(16nk^3rd)}{2k} + \frac{\log(rk)}{r} + H_2(5\delta)/2 + \delta \log d.$$

Proof. Any singleton-free $(k-1)$-hike defines a connected graph $\mathcal{H}$ with at most $k-1$ edges and, therefore, at most $k-1$ vertices. If there is no backtracking step, then all vertices except the start have a degree of at least two. Otherwise, the endpoint of one of the backtracking steps may have a degree of 1. Thus, there are at most 2 vertices of degree one. When k is unbounded, we use the bound from the first encoding i.e. Lemma 3.3.5 and combine it with the number of possible hikes on this from Claim 3.3.8 to get

$$\begin{aligned}
&\le 2n \cdot d \cdot (d-1)^{k-2} \cdot 2^{2(k-1)} (r(k-1))^{\frac{2(k-1)}{r}+1} \\
&\le (nrk) \cdot (d-1)^k \cdot 2^{2k} (rk)^{\frac{2k}{r}} \\
&\le \left(2 \cdot 2^{\log(nrk)/2k} 2^{\frac{\log(rk)}{r}}\right)^{2k} (d-1)^k \\
&\le \left(2^{\gamma_1}\sqrt{d-1}\right)^{2k}.
\end{aligned}$$

The assumption on k lets us use Corollary 3.3.10 which when combined with Lemma 3.3.5 gives us the bound on the number of such graphs as $4nk^2d \cdot (d-1)^{k-1} \cdot \binom{k}{2\delta k+1} \cdot d^{2\delta k+1}$. Combining with the number of possible hikes on this from Claim 3.3.8, we get the total number of singleton-free k-hikes bounded by

$$\begin{aligned}
&\le 4n(k-1)^2 \cdot d \cdot (d-1)^{k-2} \cdot \binom{k-1}{2\delta(k-1)+2} \cdot d^{2\delta(k-1)+2} (r(k-1))^{\frac{2k-2}{r}+1} \\
&\le (16nk^3rd)(d-1)^k \cdot 2^{H_2(5\delta)k} \cdot d^{2\delta k} (rk)^{\frac{2k}{r}} \\
&\le \left(2^{\log(16nk^3rd))/2k} d^{\delta} 2^{\frac{\log(rk)}{r}} 2^{H_2(5\delta)/2}\right)^{2k} (d-1)^k \\
&\le \left(2^{\gamma_2}\sqrt{d-1}\right)^{2k}.
\end{aligned}$$

□

3.4 How we use this method

In this section, we will use the bound on singleton-free hikes obtained in the last section to bound the eigenvalue of the lifted graph. We first handle non-singleton-free hikes and show that they can be easily bounded by the ε-biased property of the distribution of the signings. We then formalize the construction by instantiating it using an expander from MOP having large bicycle-free radius and then bring the bounds together.

3.4.1 A Simple Generalization of The Trace Power Method

We now show that the problem of bounding the spectral radius of the signed non-backtracking operator reduces to counting singleton-free hikes. This reduction is a straightforward generalization of the argument [MOP20, Prop. 3.3] for $\mathbb{Z}_2$ to any abelian group.

Let $B_s(\chi)$ (as defined in Definition 3.1.1) be the signed non-backtracking operator with respect to a signing and a non-trivial character χ and $\rho(B_s)$ denote its spectral radius. The goal is to bound the largest eigenvalue of $B_s(\chi)$. The trace method is the name for utilizing the following inequality,

$$\rho(B)^{2k} \leq \sum_i |\lambda_i^k|^2 = \mathrm{tr}((B^*)^k B^k) = \|B^k\|_F.$$

The signing s is drawn from some distribution $\mathcal{D}$, and we wish to show via the probabilistic method that there exists a signing in $\mathcal{D}$ for which $\rho(B_s(\chi))$ is small for any set of $(l-1)$ non-trivial characters χ. We will use a first-order Markov argument and therefore wish to bound $\mathbb{E}_{s\sim\mathcal{D}}[\mathrm{tr}(B_s^k(B_s^*)^k)]$. Writing it out, we get,

$$\begin{aligned} T_\chi(s) = \mathrm{tr}((B_s^*)^k B_s^k) &= \sum_{e\in E^d} \left((B_s^*)^k B_s^k e\right)_e \\ &= \sum_{(e_0,\cdots,e_{2k})} B(e_0,e_1)\cdots B(e_{k-1},e_k) B^*(e_k,e_{k+1})\cdots B^*(e_{2k-1},e_{2k}) \end{aligned}$$

$$= \sum_{(e_0,\cdots,e_{2k})} \chi(s(e_1))\cdots\chi(s(e_k))\chi^*(s(e_k))\cdots\chi^*(s(e_{2k-1}))$$
$$= \sum_{(e_0,\cdots,e_{2k})} \chi(s(e_1))\cdots\chi(s(e_{k-1}))\chi^*(s(e_{k+1}))\cdots\chi^*(s(e_{2k-1})).$$

Notice that e_0, e_k do not appear in the term, and so we define $\mathcal{H}_{k-1}$ as the multiset of all tuples $(e_1, \ldots, e_{k-1}, e_{k+1}, \ldots, e_{2k-1})$ appearing in the support of this summation. We denote each term in the summation above by $\chi_w(s)$ where $w \in \mathcal{H}_{k-1}$. It follows directly from the definition that each $w \in \mathcal{H}_{k-1}$ defines a $(k-1)$-hike. Also, observe that any tuple appears at most $(d-1)^2$ times as given a tuple w, we have at most $(d-1)$ choices for each e_0, e_k. Let $\mathcal{H}^s_{k-1}$ denote the singleton-free hikes in $\mathcal{H}_{k-1}$. We can split $T_\chi(s) = T_1(s) + T_2(s)$ where

$$T_1(s) = \sum_{w\in\mathcal{H}^s_{k-1}} \chi_w(s), \quad T_2(s) = \sum_{w\notin\mathcal{H}^s_{k-1}} \chi_w(s).$$

We now recall the definition of ε-biased distributions (Definition 1.3.10) specialized to Abelian groups. This will be the key pseudorandomness tool.

Definition 3.4.1 (Bias)**.** Let $\mathcal{D}$ be a distribution on a group G. The distribution, $\mathcal{D}$, is ν-biased if $\mathrm{bias}_\chi(\mathcal{D}) := |\mathbb{E}_{h\sim\mathcal{D}}\chi(h)| \le \nu$ for every non-trivial character, χ.

Lemma 3.4.2. *Let $\mathcal{D} \subseteq G^E$ be an ν-biased distribution and let $w \notin \mathcal{H}^s_{k-1}$ be a singleton-hike, i.e., there is an edge that is traveled exactly once. Then, $|\mathbb{E}_{s\sim\mathcal{D}}\chi_w(s)| \le \nu$.*

Proof. Let the set of distinct edges in w be $\{e_1, \cdots, e_r\}$ and let edge e_i be traveled t_i times where t_i takes the sign into account.[5]Let e_j be the edge traversed exactly once. Then, $t_j = \pm 1$. Now, we can rewrite $\chi_w(s) = \prod_{i=1}^r \chi(s(e_i))^{t_i}$ and it can be extended to a character on $H^{E(G)}$. Since $t_j = \pm 1$, this character is non-trivial, and the claim follows from the ν-biased property. □

5. Let e_i appear f_1 times in the first $k-1$ steps and b_1 times in the next $(k-1)$ steps. Similarly let e_i^T which is the reverse direction of e appear f_2 times in the first $k-1$ steps and b_2 times in the next $(k-1)$ steps. Then, $t_i = f_1 + b_2 - f_2 - b_1$.

Lemma 3.4.3 (Analog of Corr. 3.11 in [MOP20])**.** *Let* X *be a* d*-regular graph on* n*-vertices,* $\varepsilon < 1$ *be a fixed constant, and* G *be an abelian group. Let* $\ell := |G|$*, and let*$\mathcal{D} \subseteq G^m$ *be an* ν*-biased distribution such that* $\nu \leq (n\ell d^2)^{-1} \cdot \left(\frac{\varepsilon}{d}\right)^{2k}$.

Assume that the number of singleton-free $(k-1)$*-hikes is bounded by* $(2^\gamma\sqrt{d-1})^{2k}$*. Then for any non-trivial character* χ *of* G*, we have that except with probability at most* $1/\ell$ *over* $\mathcal{D}$*,* $\rho(B(\chi)) \leq 2^{\gamma'}\sqrt{d-1} + \varepsilon$ *where* $\gamma' = \gamma + \frac{\log(\ell d^2)}{2k}$.

Proof. By the decomposition above, we have $T(s) = T_1(s) + T_2(s)$. As each term in the expression is of the form $\chi(h)$ and as remarked earlier, all the characters are roots of unity so $|\chi(s(e))| = 1$. Thus, $|T_1(s)| \leq |\pi^{-1}(\mathcal{H}^*_{k-1})| \leq (d-1)^2\,|\mathcal{H}^*_{k-1}|$.

$$\begin{aligned}
\mu := |\mathbb{E}_{s\sim\mathcal{D}} T| &= |\mathbb{E}T_1 + \mathbb{E}T_2| \\
&\leq |\mathbb{E}T_1| + |\mathbb{E}T_2| \\
&\leq |\mathcal{H}^s_{k-1}| + \sum_{w\notin\mathcal{H}^s_{k-1}} |\mathbb{E}_{s\sim\mathcal{D}}\chi_w(s)| \\
&\leq d^2(2^\gamma\sqrt{d-1})^{2k} + \nu|\mathcal{H}_{k-1}| \\
&\leq d^2(2^\gamma\sqrt{d-1})^{2k} + \nu n d^{2k+2}.
\end{aligned}$$

Here we have used the observation that $|\mathcal{H}^s_{k-1}| \leq (d-1)^2\{|\text{Singleton-free } (k-1)\text{-hikes}|\}$ and Lemma 3.4.2. The bound on $|\mathcal{H}_{k-1}|$ is trivial as we have nd choices for the starting edge and a walk of length of $2k+1$. Since T is a non-negative random variable, we apply Markov to conclude that $T \leq \mu\ell$ with probability at most $1/\ell$. Hence, we obtain,

$$\begin{aligned}
\rho(B_s(\chi)) \leq T^{1/2k} < (\mu\ell)^{1/2k} &\leq \left(d^2\ell\left(2^\gamma\sqrt{d-1}\right)^{2k} + \nu\ell n d^{2k+2}\right)^{1/2k} \\
&\leq (d^2\ell)^{1/2k}2^\gamma\sqrt{d-1} + \left(\nu\ell n d^{2k+2}\right)^{1/2k} \\
&\leq 2^{\gamma'}\sqrt{d-1} + (\nu\ell n d^2)^{1/2k}d \\
&\leq 2^{\gamma'}\sqrt{d-1} + \frac{\varepsilon}{d}d \leq 2^{\gamma'}\sqrt{d-1} + \varepsilon.
\end{aligned}$$

□

3.4.2 The Instantiation

We will need two tools before we instantiate the explicit construction of abelian lifted expanders leading to Theorem 3.0.1. The first is an explicit construction of expander graphs to be used as base graphs in the lifting operation. Since we need this technical condition of bicycle-freeness, we use the construction in [MOP20].

Theorem 3.4.4 ([MOP20, Theorem 1.1])**.** *For any given constants* $d \geq 3, \varepsilon > 0$*, one can construct in deterministic polynomial time, an infinite family of graphs* $\{X_n\}$ *with* $\lambda(X_n) \leq 2\sqrt{d-1} + \varepsilon$ *and* X_n *is*

- $n \leq |V(X_n)| \leq 2n$,
- X_n *is bicycle-free at radius* $c \log_{d-1}(|V(X_n)|)$,
- $\lambda_2(B_X) \leq \sqrt{d-1} + \varepsilon$.

The second tool is a ν-biased distribution for abelian groups (having a sample space depending polynomially on $1/\nu$). We use a recent construction by Jalan and Moshkovitz.

Theorem 3.4.5 ([JM21])**.** *Given the generating elements of a finite abelian group* G *and an integer* $m \geq 1$ *and* $\nu > 0$*, there is a deterministic polynomial time algorithm that constructs subset* $S \subseteq G^m$ *with size* $O\left(\frac{m \log(H)^{O(1)}}{\nu^{2+o(1)}}\right)$ *such that the uniform distribution over* S *is* ν*-biased.*

We are now ready to prove our main result.

Theorem 3.0.1. *For large enough* n *and constant degree* $d \geq 3$*, given an Abelian group* G *such that* $|G| \leq \exp(n^{\Theta(1)})$*, and any fixed constant* $\varepsilon \in (0, 1)$*, we can construct in deterministic polynomial time, a* d*-regular graph* X *on* $\Theta(n\ell)$ *vertices such that,*

- X *is* G*-lift of a graph* X_0 *on* $\Theta(n)$ *vertices.*
- *If* $|G| \leq \exp\left(n^{\delta(d,\varepsilon)}\right)$*, then* $\lambda(X) \leq 2\sqrt{d-1} + \varepsilon$.

- *If* $|G| \leq \exp\left(n^{\delta}\right)$ *and also* $d \geq d_0(\varepsilon)$, *then* $\lambda(X) \leq \varepsilon \cdot d$.

Proof. Construct X_0 on $n \leq n' \leq 2n$ vertices for given (d, ε) using Theorem 3.4.4 which has $r \geq c \log_{d-1} n'$. Let $\ell := |G|$.

- **Regime 1** - Here, shorter walks will suffice, and we will use the bound on γ_2 from Lemma 3.3.11. To get near-Ramanujan, we need $\gamma' = \gamma_2 + \frac{\log(d^2\ell)}{2k} = \gamma_2' + \frac{\log(\ell)}{2k}$ to be vanishing with ε. Observe that when $k = \omega(\log n)$, γ_2 is bounded by $o(1) + \left(2\sqrt{\delta} + \delta \log d\right)$. We pick δ small enough and assume that $n' \geq N(\varepsilon, d)$ such that $\gamma_2' \leq \frac{2\varepsilon}{\sqrt{d-1}}$. In the bounded k regime, we can pick $k < e^{\delta r}$. Since, $\frac{\log(\ell)}{2k}$ must also be vanishing in ε, this forces $\log(\ell) \leq \varepsilon k \leq \varepsilon e^{\delta r}$. This explains the bound on ℓ.

- **Regime 2** - Here ℓ is larger and so we pick $k = \log \ell$. Now, we need to use γ_1 which we recall is $1 + \frac{\log k}{r} + o(1)$. Thus, $\gamma' = (\gamma_1 + \frac{\log d^2}{k}) + \frac{\log \ell}{k} \leq 3/2 + \frac{\log k}{r}$. Since, $r = c \log_{d-1}(n')$, to get non-trivial expansion $k \leq n^{c/2}$ which explains the bound on the exponent δ.

The precise parameters are as follows:

Regime	δ	k	ν	γ'
1	$O\left(\frac{\varepsilon^2}{d}\right)$	$\frac{10\sqrt{d-1}}{\varepsilon} \max(\log \ell, \log n)$	$(n\ell d^2)^{-1}\left(\frac{\varepsilon}{3d}\right)^{2k} = (n\ell)^{c_{d,\varepsilon}}$	$\frac{2\varepsilon}{3\sqrt{(d-1)}}$
2	$\leq c/2$	$\log \ell = n^{\delta}$	$(n\ell d^2)^{-1}\left(\frac{1}{3d}\right)^{2k} = (n\ell)^{c_d}$	$2 + \frac{\delta}{c}\log(d-1)$

Table 3.1: Precise parameters for the different regimes

Construct a ν-biased distribution $\mathcal{D}$ using Theorem 3.4.5. These two constructions take $\mathrm{poly}(n, \ell)$ time. From Corollary A.2.2, we have to analyze $B(\chi)$ for $\ell - 1$ non-trivial characters χ that appear in this decomposition. The largest eigenvalue is given by $B(1)$, which is $d - 1$. For the second largest, $\lambda_2(B(1)) \leq \sqrt{d-1} + \varepsilon$ by the property of the base graph G obtained by Theorem 3.4.4. Since we have the bicycle-free property, we can use Lemma 3.4.3 to

conclude that for any non-trivial character, the following hold with probability at least $1 - 1/\ell$,

- **Regime 1** - $\rho(B(\chi)) \leq 2\sqrt{\gamma'}\sqrt{d-1} + \varepsilon/3 \leq \sqrt{d-1} + \varepsilon$.
- **Regime 2** - $\rho(B(\chi)) \leq 2\sqrt{\gamma'}\sqrt{d-1} + 1 \leq 2 \cdot 2^2 d^{\delta/c}\sqrt{d-1} \leq \varepsilon d$ when $d \geq \left(\frac{8}{\varepsilon}\right)^{\frac{2c}{c-2\delta}}$.

Using Fact 3.1.3, we assume that the decomposition has exactly one trivial character (say, χ_1) and $(l-1)$ non-trivial characters. Then, for the trivial character $\rho(B_{X_0(s)}) = \rho(B(\chi_1)) = d - 1$ and thus, $\rho_2(B) = \max\left\{\lambda(X_0), \max_{i=2}^{\ell} \rho(B(\chi_i))\right\}$.

Since the bound holds for any non-trivial χ except with probability $1/\ell$, we take a union bound over these $\ell - 1$ characters, we get that there is a labeling $s \in D$ such that the bound holds for $\rho(B(\chi_i))$ and thus for $\lambda(B_{X_0(s)})$. By Fact 3.1.2, we get that $\lambda(G) \leq 2\rho_2(B_X)$ which satisfies the bounds we need.

We can brute force through each $s \in \mathrm{Supp}(\mathcal{D})$ to find an s such that the lifted graph $X = X_0(s)$ has the required spectral gap. Checking this is a simple linear algebraic task and can be done in time cubic in $n\ell$. Therefore, the total time taken is $\mathrm{poly}(n, \ell)$. □

CHAPTER 4

DERANDOMIZED POWERING

Piro: Ready to brace the polar vortex next week?

Laqa: Yes, I have saved the papers on my Zotero! What paper are you holding, Janki?

Janki: This? It is Ta-Shma's paper on binary codes near the GV bound.

Piro: Another way to say it is that starting with an ε_0-expander, $\mathrm{Cay}(\mathbb{Z}_2^n, S)$, they amplify it to an ε-expander $\mathrm{Cay}(\mathbb{Z}_2^n, S')$ of degree roughly $\frac{\log n}{\varepsilon^{2+o(1)}}$.

Laqa: So, does this technique work for other groups? What if I start with a constant degree Cayley expander over a non-abelian group? Where is $\mathbb{Z}_2^n$ used?

Janki: The construction is very combinatorial; just walks on a graph. The analysis is via a *bias operator*, that captures the bias of S. This is ± 1-valued for $\mathbb{Z}_2^n$, but I do not know what it would look like for general groups.

Piro: Also, the degree bound achieved via this is what you get randomly via Alon–Roichmann. But what you are saying would imply a constant degree expander, which has no randomized proof.

Laqa: That makes sense. However, I am starting with a strong object that beats the random argument, so it may not be hopeless.

A different paradigm to explicitly construct almost optimal expanders is to start with a family of weak expanders and amplify it to a strong one via combinatorially defined graph products. Here, weak versus strong refers to the trade-off between degree and expansion. Alon–Bopanna bound says that to achieve spectral expansion λ, the degree of

the graph must be large, $d \geq \Omega(1/\lambda^2)$. Graphs that achieve the optimal degree are called *Ramanujan graphs*. We will call any graph that achieves $d \leq 1/\lambda^{2+o(1)}$ as near-optimal or almost Ramanujan.

Graph powering is a simple way to improve expansion that transforms a graph X to X^t, where X^t has A_X^t as its adjacency matrix[1]. If X was a λ_0-expander, then X^t has expansion λ_0^t. However, the degree also increases proportionally, $d' = d_0^t$, and the trade-off remains the same. Hence, one needs a "derandomized power," i.e., a low-degree subgraph of X^t that retains its expansion. Since we wish to build graphs with symmetry, we will require that if $X = \mathrm{Cay}(G, S)$ is a Cayley graph, then X' must also be a Cayley graph. In other words, we want to efficiently compute a sparse subset of S^t that retains the expansion.

We now have a question about subsets of a group G, and the key property we need to analyze is the *bias* of this set, as $\mathrm{Cay}(G, S)$ is an ε-expander if and only if S is an ε-biased set. Recall that a set $S \subseteq G$ is called an *ε-biased set* if for every non-trivial irreducible *representation*, ρ, of G, we have $\|\mathbb{E}_{s\sim S}[\rho(s)]\|_{\mathrm{op}} \leq \varepsilon$. Since this notion of bias is important for us, we first rephrase the technical result of our paper from the perspective of "bias amplification".

4.1 Bias Amplification

Let $f : S \to M_\ell(\mathbb{C})$ be a matrix-valued function. The quantity $\|\mathbb{E}_{s\sim S}[f(s)]\|_{\mathrm{op}}$ is known as the *bias* of the operator-valued function f with respect to S. The key idea in the prior works (and ours) is to establish an "bias amplification" result of the following form:

Theorem 4.1.1 (Template Amplification Result)**.** *Let S be a finite set and $\lambda_0 \in (0, 1)$ be a constant. For every $\lambda > 0$, there exists a deterministic polynomial time algorithm to construct $\mathcal{W} \subseteq S^t$ of size $|\mathcal{W}| \leq \mathrm{poly}(|S|, 1/\lambda)$ such that*

1. The more standard definition does not include multiplicities, i.e., X^t has an edge between vertices of distance at most t.

- **Scalar Amplification** *For every function* $f : S \to \mathbb{C}$ *such that* $\|f\|_\infty \le 1$, *if* $|\mathbb{E}_{s\sim S}[f(s)]| \le \lambda_0$ *then we have* $|\mathbb{E}_{(s_1,\cdots,s_t)\sim W}[f(s_t)\cdots f(s_1)]| \le \lambda$.
- **Operator Amplification** *For every function* $f : S \to M_\ell(\mathbb{C})$ *such that* $\max_s \|f(s)\|_{\mathrm{op}} \le 1$, *if* $\|\mathbb{E}_{s\sim S}[f(s)]\|_{\mathrm{op}} \le \lambda_0$ *then we have* $\left\|\mathbb{E}_{(s_1,\cdots,s_t)\sim W}[f(s_t)\cdots f(s_1)]\right\|_{\mathrm{op}} \le \lambda$.

Corollary 4.1.2 (Amplification of Cayley Graphs)**.** *If* $\mathrm{Cay}(G, S)$ *is a given* λ_0*-expander, and for any* $\lambda > 0$*, the set* W *be the output of Theorem 4.1.1. Set* $S' = \{s_t \cdots s_1 \mid (s_1, \cdots, s_t) \in W\}$. *Then,* $\mathrm{Cay}(G, S')$ *is a* λ*-expander of degree* $|W|$.

Proof. Apply Theorem 4.1.1 for each irreducible representation ρ. □

Note that over Abelian groups, the scalar amplification suffices as the irreducible representations of Abelian groups are 1-dimensional, i.e., scalar-valued functions called *characters*. However, for general finite groups, one needs the amplification result for operator-valued functions as the irreps are matrix-valued (of dimensions up to $\sqrt{|G|}$).

Random Amplification is unknown! A first attempt would be to use matrix Chernoff plus union bound to select a random subset S' such that $\|\mathbb{E}_{s\in S'}[\rho(s)]\|_{\mathrm{op}}$ is small for any irreducible representation ρ. This works in the scalar case, but in general, this approach requires $|S'| \ge \Omega(\log \dim(\rho))$. For non-abelian groups, we have irreducible representations such that $\dim(\rho) = \mathrm{poly}(|G|)$; therefore, this cannot deduce the existence of constant-sized subsets that achieve expansion. This difficulty is also present in the proof of the Alon–Roichman theorem [AR94] and the reason why the only known generic upper bound for non-Abelian groups uses $\Omega(\log(|G|))$ random generators to obtain an expander.

Prior Results on Bias Amplification

Scalar amplification Much of the earlier work has focused on the case of Abelian groups, especially $G = \mathbb{Z}_2^n$, as this has connections to other objects like error-correcting codes.

Rozenman and Wigderson introduced the use of expanders for "scalar amplification" via an iterated application of the expander mixing lemma. Alon (in an unpublished email[2]) introduced the idea of using walks on an (auxiliary) expander graph X, whose vertices are identified with elements of S. The set $\mathcal{W} \subseteq S^t$ is chosen to be the collection of all walks of length $(t-1)$ on X. This technique gives a λ-biased set $\mathcal{W}$, of size $|\mathcal{W}| \leq O(|S|/\lambda^{4+o(1)})$ (*cf.*, [TS17]), which is quite good but still sub-optimal.

Ta-Shma [TS17] managed to close the gap almost optimally ($\lambda^{-2-o(1)}$) using the *s-wide replacement product* to derandomize the above amplification. The s-wide replacement product of Ben-Aroya and Ta-Shma [BATS08] is a higher-order version of the zig-zag product [RVW00]. Using the collection of walks on the s-wide replacement product allows for a much smaller collection $\mathcal{W} \subseteq S^t$ with nearly optimal size. This scalar technique was later applied to the more general case of arbitrary Abelian groups by Jalan and Moshkowitz [JM21].

Operator amplification To extend Ta-Shma's approach to non-Abelian groups, it is necessary to work with operator-valued functions, $f : S \to M_\ell(\mathbb{C})$, as the irreducible representations are no longer of dimension one. To the best of our knowledge, only one general result was known for general groups. Chen, Moore, and Russell [CMR13] analyzed the above expander walk construction using a matrix version of the expander mixing lemma. This gives an amplification procedure for Cayley graphs of general groups, but the resulting degree $O(|S|/\lambda^{11})$ to achieve final expansion λ is sub-optimal.

4.1.1 Overview of Our Results

The main contribution of this work is the identification of appropriate natural linear algebraic extensions to Ta-Shma's amplification framework [TS17]. This gives an almost-

2. We thank the anonymous reviewer for pointing out this reference.

optimal *dimension independent* generalization of the scalar amplification result to operator-valued functions. Our result sharpens that of Chen, Moore and Russell [CMR13] by reducing the degree from $O(|S|/\lambda^{11})$ to $O(|S|/\lambda^{2+o(1)})$

Theorem 4.1.3 (Operator Amplification (this work))**.** *Let S be a finite set and $\lambda_0 \in (0,1)$ be a constant. For every $\lambda > 0$, there exists a deterministic polynomial time algorithm to construct $W \subseteq S^t$ of size $|W| \leq O(|S|/\lambda^{2+o(1)})$ such that for every function $f : S \to M_\ell(\mathbb{C})$ with $\|\mathbb{E}_{s\sim S}[f(s)]\|_{op} \leq \lambda_0$ and $\max_s \|f(s)\|_{op} \leq 1$, we have $\|\mathbb{E}_{w\sim W}[f(w)]\|_{op} \leq \lambda$.*

The key extension is a simple and yet extremely useful change in the bias operator (Π_f) defined by Ta-Shma, which is a central object in the analysis of both [TS17] and [JM21]. In both these cases, f is scalar, and they define,

$$\Pi_f : \mathbb{C}[S] \to \mathbb{C}[S] \text{ where } \Pi_f \cdot s = f(s) \cdot s\,.$$

However, this approach is not readily generalizable to operators, and the view we take is that if $f : S \to M_\ell(\mathbb{C})$, then Π_f is actually an operator on $\mathbb{C}^\ell \otimes \mathbb{C}[S]$ defined as.

$$\Pi_f : \mathbb{C}^\ell \otimes \mathbb{C}[S] \to \mathbb{C}^\ell \otimes \mathbb{C}[S] \text{ where } \Pi_f\,(v \otimes s) = f(s)\,v \otimes s\,.$$

In the Abelian case, we have $\ell = 1$ and this is isomorphic to the setup by Ta-Shma. This generalization is very natural, and we show that not only does the older machinery gel well with this, but the proof remains intuitive with the different spaces neatly delineated.

More precisely, we first establish an *operator version* of the expander walk amplification, and then we derandomize it using (a suitable version of) the s-wide replacement product. Furthermore, since the result does not depend on the dimension, ℓ, we can use it even for functions $f : S \to \mathcal{L}(\mathcal{H})$ where $\mathcal{L}(\mathcal{H})$ is the space of bounded linear operators on an arbitrary Hilbert space, $\mathcal{H}$, possibly infinite dimensional. This is useful if the underlying group is not finite but finitely generated by S.

4.1.2 Notation

Since we deal with various vector spaces and graphs, we will find it useful to establish some convenient notation for this chapter. The following is a summary for ready reference.

- The main multigraphs we study will be X and Y with vertices V_X, V_Y and normalized adjacency operators A_X, A_Y.
- We denote vertices of X, Y by x, y and an ordered tuple of vertices by $\vec{x} = (x_0, \cdots, x_t)$.
- We use u, v, w to denote arbitrary vectors in $\mathcal{H}$ and x, y for basis vectors of $\mathbb{C}[V_X], \mathbb{C}[V_Y]$ where $\mathbb{C}[V_X]$ is the complex vector space with the elements of V_X being an orthonormal basis.
- The tensored vector spaces have an induced inner product. For $\mathcal{X}_\mathcal{H} \coloneqq \mathcal{H} \otimes \mathbb{C}[V_X]$, it is $\langle v \otimes x, w \otimes x' \rangle = \langle v, w \rangle_\mathcal{H} \langle x, x' \rangle$. Similarly, we have one on $\mathcal{XY}_\mathcal{H} \coloneqq \mathcal{X}_\mathcal{H} \otimes \mathbb{C}[V_Y]$.
- Orthogonal decomposition: $\mathcal{X}_\mathcal{H} = \mathcal{X}^{\parallel}_\mathcal{H} \oplus \mathcal{X}^{\perp}_\mathcal{H}$ where $\mathcal{X}^{\parallel}_\mathcal{H} \coloneqq \operatorname{span}\{v \otimes \vec{1} \mid v \in \mathcal{H}\}$. Here, $\vec{1}$ denotes the un-normalized all-ones vector. Similarly, $\mathcal{XY}_\mathcal{H} = \mathcal{XY}^{\parallel}_\mathcal{H} \oplus \mathcal{XY}^{\perp}_\mathcal{H}$, where $\mathcal{XY}^{\parallel}_\mathcal{H} \coloneqq \operatorname{span}\{z \otimes \vec{1} \mid z \in \mathcal{X}_\mathcal{H}\}$.
- The operator $\mathring{A}$ denotes the extension of operator A to a tensor product of spaces where it acts as identity on the other spaces. For example, A_X acts on $\mathbb{C}[V_X]$ and its extension to $\mathcal{X}_\mathcal{H}$ is $\mathring{A}_X = I_\mathcal{H} \otimes A_X$. However, if we were working on $\mathcal{XY}_\mathcal{H}$, it would be $\mathring{A}_X = I_\mathcal{H} \otimes A_X \otimes I_Y$ instead[3].
- Given an operator-valued function $f : V_X \to \mathcal{L}(\mathcal{H})$, the generalized *bias operator* is defined on the basis as[4],

$$\Pi_f : \mathcal{X}_\mathcal{H} \to \mathcal{X}_\mathcal{H}, \quad v \otimes x \mapsto f(x)\, v \otimes x.$$

3. The spaces will be self-evident and the use of the same notation should not be confusing.

4. An equivalent matrix definition is $\Pi_f \coloneqq \sum_{x \in V_X} f(x) \otimes E_{x,x}$ where $E_{x,x} \in \mathbb{C}^{V_X \times V_X}$ is the diagonal matrix with exactly one non-zero entry of value 1 in the row and column indexed by the vertex x.

4.2 Derandomized Powering via Expander Walks

In this section, we establish a new *operator* analog of the expander walk–based bias amplification procedure for *scalars*. An analysis of this scalar amplification was given by Ta-Shma in [TS17]. We prove an operator analog that can amplify from any bias (Theorem 4.2.6) which implies the main result below (Theorem 4.2.1), that amplifies from constant bias.

Theorem 4.2.1 (Operator Amplification via Expander Walks)**.** *Let* X *be a* $\lambda(X)$*-spectral expander, and let* $\mathcal{W}_t$ *be the collection of walks obtained from walks of length* t *on* X*. Then for any operator valued function* f *such that* $\left\|\mathbb{E}_{x\in V_X}[f(x)]\right\|_{op} \le \lambda_0$ *and* $\max_{x\in V_X}\|f(x)\|_{op} \le 1$*, we have*

$$\left\| \mathop{\mathbb{E}}_{\vec{x}\in\mathcal{W}_t}[f(x_t)\cdots f(x_0)] \right\|_{op} \le (2\lambda(X)+\lambda_0)^{\lfloor t/2\rfloor} .$$

We remark that a precursor of these techniques, in the simpler setting of Abelian groups appears in the pioneering work of Naor and Naor introducing ε-biased distributions over the group $\mathbb{Z}_2^m$ using expanders [NN90].

This simpler amplification of Theorem 4.2.1 will be crucially used to bootstrap the almost-optimal amplification. Moreover, it yields a construction of expanding Cayley graphs close to any desired size, which will be required later.

This bias reduction procedure uses walks on an auxiliary expander graph. Here, we only use its expansion property (as opposed to later when we rely on its structure for the s-wide construction). With this, it is already possible to obtain $1/\lambda^{4+o(1)}$ dependence on the final degree of an λ-expander.

Theorem 4.2.2. *Let* $S \subseteq G$ *such that* $\lambda(\mathrm{Cay}(G,S)) = \lambda_0 < 1$*. For every* $\lambda \in (0,1)$ *and constant* $\beta \in (0,1)$*, we can find* $S' \subseteq G$ *in time* $\mathrm{poly}(|S|, 1/\lambda_0, 1/\lambda)$ *such that* $\lambda(\mathrm{Cay}(G,S')) \le \lambda$ *and* $|S'| = O_{\lambda_0}\left(\frac{|S|}{\lambda^{4+\beta}}\right)$.

4.2.1 Operator Norm Decay

Lemma 4.2.3. *Let $\mathcal{W}_t \subseteq V_X^{t+1}$ be the collection of all length t walks on the graph X and we define $\mathring{A}_X = I_{\mathcal{H}} \otimes A_X$. Then, we have*

$$\left\| \mathop{\mathbb{E}}_{\vec{x} \in \mathcal{W}_t} [f(x_t) \cdots f(x_0)] \right\|_{op} \leq \left\| \Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t \right\|_{op} \leq \left\| \left(\mathring{A}_X \, \Pi_f \right)^2 \right\|_{op}^{\lfloor t/2 \rfloor}.$$

Proof.

$$\Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t \mathop{\mathbb{E}}_{x \in V_X} [v \otimes x] = \mathop{\mathbb{E}}_{\vec{x} \in \mathcal{W}_t} [f(x_t) \cdots f(x_0) \, v \otimes x_t]. \tag{4.1}$$

This can be shown easily via induction on t, and we refer to Lemma 4.3.7 for a formal proof of a more general statement. We use projection and lifting maps to move between the spaces $\mathcal{X}_{\mathcal{H}}$ and $\mathcal{H}$. Define $P_{\mathcal{H}} : \mathcal{X}_{\mathcal{H}} \to \mathcal{H}$ and $L_{\mathcal{H}} : \mathcal{H} \to \mathcal{X}_{\mathcal{H}}$, as,

$$P_{\mathcal{H}}(w \otimes x) = w, \quad L_{\mathcal{H}}(v) = \mathop{\mathbb{E}}_{x \in V_X} [v \otimes x] = \frac{1}{|V_X|} \, v \otimes \vec{1}.$$

From the definition, $\|L_{\mathcal{H}}(v)\| = \|v\| \frac{\|\vec{1}\|}{|V_X|} = \frac{\|v\|}{\sqrt{|V_X|}}$ and thus, $\|L_{\mathcal{H}}\|_{op} = 1/\sqrt{|V_X|}$. We can use Cauchy-Schwarz to get that $\|P_{\mathcal{H}}\|_{op} = \sqrt{|V_X|}$. Now, we put this together to obtain a simple expression on the quantity we need to bound

$$\begin{aligned}
\left\| \mathop{\mathbb{E}}_{\vec{x} \in \mathcal{W}_t} [f(x_t) \cdots f(x_0)] \right\|_{op} &= \sup_{\|v\|=1} \left\| \mathop{\mathbb{E}}_{\vec{x} \in \mathcal{W}_t} [f(x_t) \cdots f(x_0)] v \right\|_2 \\
&= \sup_{\|v\|=1} \left\| P_{\mathcal{H}} \mathop{\mathbb{E}}_{\vec{x} \in \mathcal{W}_t} [f(x_t) \cdots f(x_0) \, v \otimes x_t] \right\|_2 \\
&= \sup_{\|v\|=1} \left\| P_{\mathcal{H}} \Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t \mathop{\mathbb{E}}_{x \in V_X} [v \otimes x] \right\|_2 \\
&= \sup_{\|v\|=1} \left\| P_{\mathcal{H}} \Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t L_{\mathcal{H}} \, v \right\|_2 \\
&\leq \left\| \Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t \right\|_{op} \|P_{\mathcal{H}}\|_{op} \|L_{\mathcal{H}}\|_{op} \\
&\leq \left\| \Pi_f \left(\mathring{A}_X \, \Pi_f \right)^t \right\|_{op}.
\end{aligned}$$

The last inequality follows from submultiplicativity of the operator norm and the observation that $\|\Pi_f\|_{op} = \|f\|_\infty \le 1$. □

Now that we have reduced the problem to studying the operator norm, we will study how the norm decays as we take walks. We use the decomposition, $\mathcal{X}_{\mathcal{H}} = \mathcal{X}_{\mathcal{H}}^{\|} \oplus \mathcal{X}_{\mathcal{H}}^{\perp}$ where $\mathcal{X}_{\mathcal{H}}^{\|} := \mathrm{span}\{v \otimes \vec{1} \mid v \in \mathcal{H}\}$. The decay comes from two sources. For $z \in \mathcal{X}_{\mathcal{H}}^{\perp}$, we get a decay by $\lambda(X)$ by the definition of X being an expander. Claim 4.2.4 shows that for $z \in \mathcal{X}_{\mathcal{H}}^{\|}$, we get a decay from Π_f, equal to the initial bias. We put this together in Theorem 4.2.1 to obtain the desired exponential decay.

Claim 4.2.4. *For $z \in \mathcal{X}_{\mathcal{H}}^{\|}$, we have,*

$$\left\|(\Pi_f z)^{\|}\right\|_2 \le \left\|\mathop{\mathbb{E}}_{x \in V_X}[f(x)]\right\|_{op} \cdot \|z\|_2 .$$

Proof. From definition of $\mathcal{X}_{\mathcal{H}}^{\|}$, we can assume that $z = u \otimes \vec{1}$. Computing we have,

$$\begin{aligned}
\left\|\left(\Pi_f\left(u \otimes \vec{1}\right)\right)^{\|}\right\|_2 &= \sup_{w \in \mathcal{H}:\, \|w \otimes \vec{1}\|_2 = 1} \left|\left\langle w \otimes \vec{1}, \Pi_f(u \otimes \vec{1})\right\rangle\right| \\
&= \sup_{w \in \mathcal{H}:\, \|w \otimes \vec{1}\|_2 = 1} \left|\left\langle w \otimes \vec{1}, \Pi_f\left(u \otimes \sum_{x \in V_X} x\right)\right\rangle\right| \\
&= \sup_{w \in \mathcal{H}:\, \|w \otimes \vec{1}\|_2 = 1} \left|\left\langle w \otimes \vec{1}, \sum_{x \in V_X} (f(x)\, u \otimes x)\right\rangle\right| \\
&= \sup_{w \in \mathcal{H}:\, \|w \otimes \vec{1}\|_2 = 1} \left|\sum_{x \in V_X} \langle w, f(x)\, u\rangle \left\langle \vec{1}, x\right\rangle\right| \\
&= \sup_{w \in \mathcal{H}:\, \|w \otimes \vec{1}\|_2 = 1} \left|\left\langle w, |V_X|\left(\mathop{\mathbb{E}}_{x \in V_X}[f(x)]\right) u\right\rangle\right| \\
&\le \left\|\mathop{\mathbb{E}}_{x \in V_X}[f(x)]\right\|_{op} |V_X| \|w\| \|u\| = \left\|\mathop{\mathbb{E}}_{x \in V_X}[f(x)]\right\|_{op} \|z\|_2 .
\end{aligned}$$

□

The last line follows as $\|z\|_2 = \|u\|_2 \sqrt{|V_x|}$ and $\|w\| = \frac{1}{\sqrt{|V_x|}}$.

We show that for every two[5] steps of the walk, the norm of the (associated) operator decays as follows.

Lemma 4.2.5. *Let X be a $\lambda(X)$-spectral expander and let f be such that $\left\|\mathbb{E}_{x\in V_X}[f(x)\,]\right\|_{op} \le \lambda_0$ and $\max_{x\in V_X}\|f(x)\,\|_{op} \le 1$. Then,*

$$\left\|\left(\mathring{A}_X\,\Pi_f\right)^2\right\|_{op} \;\le\; 1-(1-\lambda(X))^2(1-\lambda_0)\,.$$

Proof. Let $A_J = J/|V(X)|$, where J is the $|V(X)|\times|V(X)|$ all ones matrix. We can write $A_X = (1-\lambda)A_J + \lambda E$, where $\lambda = \lambda(X)$ and $\|E\|_{op}\le 1$. Then

$$\begin{aligned}\left\|\mathring{A}_X\,\Pi_f\,\mathring{A}_X\right\|_{op} \;\le\; (1-\lambda)^2\left\|\mathring{A}_J\,\Pi_f\,\mathring{A}_J\right\|_{op} \;+\; \lambda(1-\lambda)\left\|\mathring{E}\,\Pi_f\,\mathring{A}_J\right\|_{op}\\ +\;(1-\lambda)\lambda\left\|\mathring{A}_J\,\Pi_f\,\mathring{E}\right\|_{op} \;+\; \lambda^2\left\|\mathring{E}\,\Pi_f\,\mathring{E}\right\|_{op}\,.\end{aligned}$$

To analyze $\left\|\mathring{A}_J\,\Pi_f\,\mathring{A}_J\right\|_{op}$, let $z\in\mathcal{X}_{\mathcal{H}}$ be a unit vector which is decomposed as $z = z^{\|}+z^{\perp}$. We have,

$$\begin{aligned}\left\|\left(\mathring{A}_J\,\Pi_f\,\mathring{A}_J\right)\left(z^{\perp}+z^{\|}\right)\right\|_2 \;&=\; \left\|\left(\mathring{A}_X\,\Pi_f\,\mathring{A}_X\right)z^{\|}\right\|_2 && (\text{As }\lambda(A_J)=0)\\ &=\; \left\|\mathring{A}_X\left(\left(\Pi_f z^{\|}\right)^{\perp}+\left(\Pi_f z^{\|}\right)^{\|}\right)\right\|_2\\ &=\; \left\|\left(\Pi_f z^{\|}\right)^{\|}\right\|_2\\ &\le\; \lambda_0. && (\text{By Claim 4.2.4})\end{aligned}$$

Thus, $\left\|\mathring{A}_J\,\Pi_f\,\mathring{A}_J\right\|_{op}\le\lambda_0$. Recall that $\|\Pi_f\|_{op}\le 1$ since $\max_x\|f(x)\,\|_{op}\le 1$, and we also have

5. This is the source of loss of a factor of 2 in the exponent (which leads to degree $O(|S|/\lambda^{4+o(1)})$ rather than the desired degree of $O(|S|/\lambda^{2+o(1)})$ we will later achieve. Note that the same loss occurs in the original zig-zag analysis of [RVW00], which was later remedied by the s-wise zig-zag of [BATS08].

$\|E\|_{\mathrm{op}}, \left\|A_J\right\|_{\mathrm{op}} \leq 1$. Then,

$$\begin{aligned}\left\|\mathring{A}_X \Pi_f \mathring{A}_X\right\|_{\mathrm{op}} &\leq (1-\lambda)^2\lambda_0 + 2\lambda(1-\lambda) + \lambda^2 \\ &= (1-\lambda)^2\lambda_0 + 1 - (1-\lambda)^2, \\ &= 1 - (1-\lambda)^2(1-\lambda_0)\,. \end{aligned}$$

□

The amplification guarantee of Theorem 4.2.1 trivializes if $2\lambda(X) + \lambda_0 \geq 1$. Nonetheless, we now show that amplification does occur under much weaker conditions, namely, whenever $\left\|\mathbb{E}_{x\in V_X}[f(x)]\right\|_{\mathrm{op}} < 1$ and the auxiliary graph X has expansion $\lambda(X) < 1$. This regime of bias amplification was instrumental in the breakthrough $\mathsf{SL} = \mathsf{L}$ result of Reingold [Rei05].

Theorem 4.2.6 (Operator Amplification via Expander Walks (strengthening of Theorem 4.2.1))**.** *Let X be a $\lambda(X)$-spectral expander and let $\mathcal{W}_t$ be the collection of walks obtained from walks of length t on X. Then for any operator valued function f such that $\left\|\mathbb{E}_{x\in V_X}[f(x)]\right\|_{\mathrm{op}} \leq \lambda_0$ and $\max_{x\in V_X}\|f(x)\|_{\mathrm{op}} \leq 1$, we have*

$$\left\| \mathop{\mathbb{E}}_{\vec{x}\in\mathcal{W}_t}[f(x_t)\cdots f(x_0)] \right\|_{\mathrm{op}} \leq \left[1 - (1-\lambda(X))^2(1-\lambda_0)\right]^{\lfloor t/2\rfloor}\,.$$

Proof. Follows by combining Lemma 4.2.3 and Lemma 4.2.5. □

This establishes that expander walks can be used to derandomize powers of an operator, itself given by an average of bounded operators, in the general case. In this derandomization, we still have an exponential norm decay, but we only "pay additional randomness" proportional to the degree of the auxiliary expander regardless of the number of operators.

4.2.2 Cayley Graphs and the Construction of Amplified Biased Sets

In this section, we will prove the following,

Theorem 4.2.2. *Let* $S \subseteq G$ *such that* $\lambda(\mathrm{Cay}(G,S)) = \lambda_0 < 1$. *For every* $\lambda \in (0,1)$ *and constant* $\beta \in (0,1)$, *we can find* $S' \subseteq G$ *in time* $\mathrm{poly}(|S|, 1/\lambda_0, 1/\lambda)$ *such that* $\lambda(\mathrm{Cay}(G,S')) \leq \lambda$ *and* $|S'| = O_{\lambda_0}\left(\frac{|S|}{\lambda^{4+\beta}}\right)$.

Towards this, we first formalize the connection between bias of a special subset of a group and the operator norm of a certain operator. The subset is obtained by taking random walks over an expander graph as mentioned above. We then proceed to bound this operator norm. Finally, we instantiate our construction with an explicit expander graph due to [Alo21].

The particular case of $S \subseteq G$ (for some group G) and the function f being a representation ρ on $\mathcal{H}$ leads to the amplification of biased sets. We will construct a new multiset $S' \subseteq G$ such if $\|\mathbb{E}_{s\sim S}[\rho(s)]\|_{\mathrm{op}} \leq \lambda_0$, then we have $\|\mathbb{E}_{s\sim S'}[\rho(s)]\|_{\mathrm{op}} \leq \lambda \ll \lambda_0$. Note here that the construction of S' is agnostic to ρ, and thus we can reduce the bias of all irreducible representations simultaneously! Assume that we have a graph X on the vertex set S. For $s \in S$, we have $f(s) = \rho(s)$ in this case. Let

$$S' \;=\; \{s_t s_{t-1} \cdots s_0 \mid (s_0, s_1, \cdots s_t) \in \mathcal{W}_t\}\,,$$

which will be our new amplified biased set. Using the homomorphism property of ρ, we have the following simplification

$$\mathop{\mathbb{E}}_{w=(s_0,\ldots s_t)\in\mathcal{W}_t}[f(s_t)\cdots f(s_0)] \;=\; \mathop{\mathbb{E}}_{(s_0,\ldots,s_t)\in\mathcal{W}_t}[\rho(s_t)\cdots\rho(s_0)] \;=\; \mathop{\mathbb{E}}_{s'\in S'}\left[\rho(s')\right], \tag{4.2}$$

$$\text{and thus, } \mathrm{bias}(S') \;\leq\; \left\|\Pi_f \left(\mathring{A}_X\, \Pi_f\right)^t\right\|_{\mathrm{op}} \tag{4.3}$$

where S' is the new biased multiset of the construction and the second inequality follows from the preceding calculation when $\mathcal{W}_t$ is a collection of walks on X.

Instantiating the Construction

To construct S', our construction requires an auxiliary expander graph X to perform walks on. One convenient source (among several) is a recent construction of Alon.

Theorem 4.2.7 (Corollary of [Alo21, Thm. 1.3])**.** *For every $\lambda \in (0,1)$, there exists a positive integer m_λ such that for every $n \in \mathbb{N}$, there is an explicit construction of a graph X on $m_\lambda n$ vertices with degree at most $9/\lambda^2$ and $\lambda(X) \le \lambda$.*

We now establish the key amplification lemma.

Lemma 4.2.8. *Let $S \subseteq G$ such that $\mathsf{bias}(S) = \lambda_0 < 1$ and let ε_0 be a constant such that $(1+2\varepsilon_0)\lambda_0 < 1$. Then, for any $\lambda > 0$, we can explicitly compute S' such that $\mathsf{bias}(S') \le \lambda$ and $|S'| = O_{\varepsilon_0\lambda_0}\left(\frac{|S|}{\lambda^{4+4\delta(\lambda_0,\varepsilon_0)}}\right)$ where $\delta(\lambda_0,\varepsilon_0) = \frac{\log(3+6\varepsilon_0)+\log(1/\varepsilon_0)}{\log(1/\lambda_0)}$.*

Proof. Pick a constant ε_0 such that $\lambda_1 := (1+2\varepsilon_0)\lambda_0 < 1$ and use Theorem 4.2.7 to obtain an explicit $(m|S|, d, \varepsilon_0\lambda_0)$-graph X. Let S_1 be the multiset consisting of m copies of S. The bias remains the same and now, $|V(X)| = |S_1|$. We construct S' by multiplying elements of t-length walks on X where $t = \lceil 2(1+\log_{\lambda_1}(\lambda))\rceil$. The size of S' is

$$\begin{aligned}
|S'| = (m|S|)\cdot d^t &= m_{\varepsilon_0\lambda_0}|S| \cdot \left(\frac{3}{\varepsilon_0\lambda_0}\right)^{4+4\log_{\lambda_1}\lambda} \\
&= m_{\varepsilon_0\lambda_0}\left(\frac{3}{\varepsilon_0\lambda_0}\right)^4 |S| \cdot \lambda^{\frac{-4\log\left(\frac{3}{\varepsilon_0\lambda_0}\right)}{\log(1/\lambda_1)}} \\
&\le O_{\varepsilon_0\lambda_0}(|S|)\cdot\lambda^{-4\left(1+\frac{\log\left(\frac{3+6\varepsilon_0}{\varepsilon_0}\right)}{\log(1/\lambda_0)}\right)}.
\end{aligned}$$

Let ρ be any irreducible representation of G. From Eq. (4.3) and Theorem 4.2.1, we get,

$$\left\| \mathop{\mathbb{E}}_{s_0\cdots s_t\in S'}[\rho(s_t\cdots s_0)] \right\|_{op} \le (2\lambda(X)+\mathsf{bias}(S))^{t/2-1} \le (\lambda_1)^{t/2-1} \le \lambda. \qquad \square$$

Using the amplification above, we now derive our first simplified explicit construction.

Theorem 4.2.2. *Let* $S \subseteq G$ *such that* $\lambda(\mathrm{Cay}(G,S)) = \lambda_0 < 1$. *For every* $\lambda \in (0,1)$ *and constant* $\beta \in (0,1)$, *we can find* $S' \subseteq G$ *in time* $\mathrm{poly}(|S|, 1/\lambda_0, 1/\lambda)$ *such that* $\lambda(\mathrm{Cay}(G,S')) \leq \lambda$ *and* $|S'| = O_{\lambda_0}\left(\frac{|S|}{\lambda^{4+\beta}}\right)$.

Proof. Pick a *constant* $\lambda' < \min\left(\frac{1}{2}, \left(\frac{3}{4}\right)^{4\beta}\right)$. This choice ensures that $\delta(\lambda', 1/2) \leq \beta$. Use Lemma 4.2.8 with target expansion λ' to obtain a set S_1 with size $|S_1| = O_{\lambda_0,\beta}(|S|)$ as λ' is a constant. Now use Lemma 4.2.8 again with S_1 as the initial set, $\varepsilon_0 = \frac{1}{2}$ and the final expansion as λ to obtain S'. Thus, the final size is $|S'| \leq O_{\lambda'}\left(\frac{|S_1|}{\lambda^{4+\delta(\lambda',1/2)}}\right) \leq O_{\lambda_0,\beta}\left(\frac{|S|}{\lambda^{4+\beta}}\right)$. □

4.2.3 Explicit Expanders close to any desired size

As an application of Theorem 4.2.2, we demonstrate an explicit construction of Cayley expanders of size $n + o(n)$ vertices for every (large enough) n. Such a construction will be crucial for us to prove Theorem 5.2.8. We cannot use existing results like the recent work of Alon [Alo21] or the construction in [TS17]. This is because Alon's construction does not have a Cayley graph structure (which our proof utilizes). On the other hand, the construction in [TS17] is a Cayley graph based on [LPS88], but it only guarantees a graph of size $O(n)$ rather than $n + o(n)$.

Recall that $SL_2(p)$ is the group of 2×2 invertible matrices over $\mathbb{F}_p$ with determinant 1. We obtain a base generating set for $SL_2(p)$ via the following result.

Theorem 4.2.9 ([Lub11])**.** *There exists an absolute constant* $\lambda_0 < 1$ *such that for every* $p > 17$, *there exists an explicit generating set* S *(of constant size independent of* p*) for* $SL_2(p)$, *such that* $\lambda\left(\mathrm{Cay}(SL_2(p), S)\right) \leq \lambda_0$.

Theorem 4.2.10 ([Che10])**.** *For every* $n \geq 2^{3\cdot 2^{15}}$, *there exists a prime in* $[n, n + 4n^{2/3}]$.

Corollary 4.2.11. *For any $n > 2^{9\cdot 2^{15}}, \lambda > 0$, there is a deterministic polynomial time algorithm to construct an (n', d, λ)-graph $\mathrm{Cay}(\mathrm{SL}_2(p), S)$, where $n' = n + O(n^{8/9})$ and $d = O(\lambda^{-4.1})$.*

Proof. Find a prime $p \in [n^{1/3} + 1, n^{1/3} + O(n^{2/9})]$, which exists due to Theorem 4.2.10, via brute-force search. Since, $\mathrm{SL}_2(p)$ is a group of order $(p^2 - 1)p$, we have $n \leq |\mathrm{SL}_2(p)| \leq n + O(n^{8/9})$. We use the constant-sized generating set S from Theorem 4.2.9 and amplify using Theorem 4.2.2. □

4.3 Derandomized Powering via the s-wide Replacement Walk

We have seen in Section 4.2 that bias reduction via random walks on an expander X is sub-optimal (by a factor of 2 in the exponent). We will derandomize this random walk construction to achieve an almost optimal bias reduction. The idea is to introduce a new graph Y, which has a much smaller degree, and to "simulate" a random walk on X via a walk on Y. This is realized by a higher-order version of the zig-zag product [RVW00] called the *s-wide replacement product* defined by Ben-Aroya and Ta-Shma [BATS08] (see Definition 4.3.5).

This section establishes our key technical result, which states that given any initial *operator valued function* of constant bias < 1, we amplify the bias in an almost optimal way.

Theorem 4.3.1 (Operator Generalization of Theorem 24 [TS17])**.** *Fix integers $t \geq s \geq 1$. Let X be any d_1-regular graph (with d_1 a power of 2), and Y be any d_2-regular Cayley graph on $\mathbb{F}_2^{s \log d_1}$. Let $s\mathcal{W}_t$ be the collection of length t walks on the s-wide replacement product of X and Y. Let $\mathcal{H}$ be a Hilbert space. For any operator valued function $f\colon V_X \to \mathcal{L}(\mathcal{H})$, satisfying $\max_{x\in V_X} \|f(x)\|_{\mathrm{op}} \leq 1$ and $\left\|\mathbb{E}_{x\in V_X}[f(x)]\right\|_{\mathrm{op}} := \lambda_0 \leq \lambda(Y)^2 - 2\lambda(X)$, we have*

$$\left\| \mathop{\mathbb{E}}_{\vec{x}\in s\mathcal{W}_t} [f(x_t)\cdots f(x_0)] \right\|_{\mathrm{op}} \leq \left(\lambda(Y)^s + s\cdot\lambda(Y)^{s-1} + s^2\cdot\lambda(Y)^{s-3}\right)^{\lfloor t/s\rfloor} \leq O_s\left(\lambda(Y)\right)^{(1-o_s(1))t}.$$

Furthermore, the size of the collection is $|s\mathcal{W}_t| = |X|\cdot d_1^s \cdot d_2^t$.

Remark 4.3.2. Note that there is an inherent trade-off between the spectral bound amplification (or the operator norm) and the degree bound (or the number of walks), which causes the suboptimality in how close this technique lets us approach the Ramanujan bound. As in [TS17], the $o_\lambda(1)$ term we obtain from the bound above is $(1/\log(1/\lambda))^c$ for some $c > 0$ (see Theorem A.1.6 for the precise computation).

Section Outline In Section 4.3.1, we recall the s-wide replacement product and describe random walks on it. Then, in Section 4.3.2, we formalize the distributions we work with and reprove the result that if Y is a Cayley graph over any product group of appropriate size, then it enables the transfer of pseudorandomness from Y to X. The key generalization to operator-valued functions is established in Lemma 4.3.7, which is identical in spirit to Eq. (4.1). In Section 4.3.3, we finish the amplification analysis similarly to [TS17]. In Appendix A.1, we provide details about instantiating the setup by explicitly constructing the graphs we need.

4.3.1 The s-wide Replacement Product and its Walks

Let X be a d_1-regular graph. For each $x \in V_X$ and $j \in [d_1]$, let $x[j]$ be the j-th neighbor of x.

Definition 4.3.3 (Locally Invertible Rotation Map)**.** X admits a locally invertible rotation map if there exists a bijection $\phi\colon [d_1] \to [d_1]$ such that for every $(x, j) \in V_X \times [d_1]$,

$$\text{if } x' = x[j], \text{ then, } x'[\phi(j)] = x\,.$$

Example 4.3.4 (Cayley Graphs are Locally Invertible)**.** Let G be a group and $A \subseteq G$ where the set A is closed under inversion. Label the neighbors of vertices in $\mathrm{Cay}(G, A)$ by elements of A such that $g[a] = a \cdot g$. Then, $\mathrm{Cay}(G, A)$ is locally invertible as the map $\phi\colon A \to A$ defined as $\phi(a) = a^{-1}$ clearly satisfies the criteria,

$$\text{if } g' = g[a] = a \cdot g, \text{ then, } g'[\phi(a)] = a^{-1} \cdot g' = g,$$

for every $g \in G$, $a \in A$.

We now define the s-wide replacement product, a generalization of the standard replacement product of graphs, which can be seen as the special case when $s = 1$.

Definition 4.3.5 (s-wide Replacement Product)**.** Suppose we are given the following:

- A d_1-regular graph X with a bijection $\phi : [d_1] \to [d_1]$ which defines a locally invertible rotation map.
- A d_2-regular graph Y on the vertex set $[d_1]^s$.

We define:

- For $i \in \{0, 1, \ldots, s-1\}$, we define $\mathrm{Rot}_i : V_X \times V_Y \to V_X \times V_Y$ such that,

$$\mathrm{Rot}_i((x, (a_0, \ldots, a_{s-1}))) := (x[a_i], (a_0, \ldots, a_{i-1}, \phi(a_i), a_{i+1}, \ldots, a_{s-1}))\,,$$

  for every $x \in V_X$ and $(a_0, \ldots, a_{s-1}) \in V_Y = [d_1]^s$. (Note that the Y component of the rotation map depends only on a vertex's Y component, not its X component.)

- Denote by X_i, the operator on $\mathbb{C}[V_X \times V_Y]$ which acts on the natural basis via the permutation Rot_i, and let A_Y be the normalized random walk operator of Y.

Then t steps of the s-wide replacement product are given by the operator

$$X_{t-1 \bmod s} \mathring{A}_Y \cdots X_{1 \bmod s} \mathring{A}_Y X_{0 \bmod s} \mathring{A}_Y \,.$$

Observe that a walk on the s-wide replacement product yields a walk on the outer graph X by recording the X component after each step of the walk. Since a walk is completely determined by its intra-cloud steps, the number of t-step walks on the s-wide replacement product is,

$$|V_X| \cdot |V_Y| \cdot d_2^t = n \cdot d_1^s \cdot d_2^t \ll n \cdot d_1^t \,,$$

which, therefore, gives us a very sparse subset of all t-walks on X. Thus, the s-wide replacement product will be used to simulate random walks on X while requiring a reduced amount of randomness (as we shall see this simulation is only possible under special conditions, namely, when we are uniformly distributed on each cloud).

4.3.2 *The Collection of Derandomized Walks*

We now describe the distribution obtained by the walks on the s-wide replacement product using the language of operators.

Definition 4.3.6 (Operators and Distributions)**.** Given a tuple of random walk operators[6] $\mathsf{B} = (\mathsf{B}_0, \cdots, \mathsf{B}_{t-1})$ on $\mathbb{C}[V_X] \otimes \mathbb{C}[V_Y]$ and a starting vertex $x_0 \in V_X$, we can define a distribution induced by the walk using these operators. More precisely, $\mathcal{D}(\mathsf{B}, x_0)$ is the distribution on $(V_X \times V_Y)^{t+1}$ such that for every $1 \le \ell \le t$,

$$(\mathsf{B}_{\ell-1} \cdots \mathsf{B}_0)\left(x_0 \otimes \frac{1}{|V_Y|}\vec{1}\right) = \mathbb{E}_{(\vec{x},\vec{y})\sim\mathcal{D}(\mathsf{B},x_0)} x_\ell \otimes y_\ell. \tag{4.4}$$

We typically suppress x_0 as it will not matter and denote $\mathcal{D}_X(\mathsf{B})$, and $\mathcal{D}_Y(\mathsf{B})$ to specify the projections to V_X, and V_Y respectively.

The next lemma is a generalization of Eq. (4.1), which we need for the s-wide replacement walk. This can also be specialized to prove Eq. (4.1) by letting Y be a graph with one vertex (and thus $\mathcal{X}_{\mathcal{H}} \cong \mathcal{X}\mathcal{Y}_{\mathcal{H}}$). Recall that $\mathring{\Pi}_{\mathsf{f}}\ (v \otimes x \otimes y) = \mathsf{f}(x)\, v \otimes x \otimes y$.

Lemma 4.3.7 (Operator Generalization)**.** *For any tuple of random walk operators* B*, any operator valued* f*, and any* $v \in \mathcal{H}$*,* $x_0 \in V_X$*, we have*

$$\left(\mathring{\mathsf{B}}_{t-1}\mathring{\Pi}_{\mathsf{f}} \cdots \mathring{\mathsf{B}}_0\mathring{\Pi}_{\mathsf{f}}\right)\left(v \otimes x_0 \otimes \frac{1}{|V_Y|}\vec{1}\right) = \mathop{\mathbb{E}}_{(\vec{x},\vec{y})\sim\mathcal{D}(\mathsf{B})}\left[\mathsf{f}(x_{t-1}) \cdots \mathsf{f}(x_0)\, v \otimes x_t \otimes y_t\right].$$

Proof. We prove the computation via induction on t. The base case is when $t = 1$.

6. Markov chain operators on $V_X \times V_Y$.

$$\left(\mathring{B}_0\mathring{\Pi}_f\right)\left(v\otimes x_0\otimes\frac{1}{|V_Y|}\vec{1}\right) = \mathring{B}_0\left(f(x_0)\,v\otimes x_0\otimes\frac{1}{|V_Y|}\vec{1}\right)$$
$$= \mathop{\mathbb{E}}_{(\vec{x},\vec{y})\sim\mathcal{D}(B)}[f(x_0)\,v\otimes x_1\otimes y_1] \quad \text{(Using Eq. (4.4) for } \ell=1)$$

Let $y_0 = \frac{1}{|V_Y|}\vec{1}$ and assume the statement holds for $t-1$. Then,

$$\begin{aligned}\left(\mathring{B}_{t-1}\mathring{\Pi}_f\cdots\mathring{B}_0\mathring{\Pi}_f\right)(v\otimes x_0\otimes y_0) &= \mathring{B}_{t-1}\mathring{\Pi}_f\cdot\prod_{i=t-2}^{0}\left(\mathring{B}_i\mathring{\Pi}_f\right)(v\otimes x_0\otimes y_0)\\ &= \mathring{B}_{t-1}\mathring{\Pi}_f\mathop{\mathbb{E}}_{(\vec{x},\vec{y})\sim\mathcal{D}(B)}[f(x_{t-2})\cdots f(x_0)\,v\otimes x_{t-1}\otimes y_{t-1}]\\ &= \mathring{B}_{t-1}\mathop{\mathbb{E}}_{(\vec{x},\vec{y})\sim\mathcal{D}(B)}[f(x_{t-1})f(x_{t-2})\cdots f(x_0)\,v\otimes x_{t-1}\otimes y_{t-1}]\\ &= \mathop{\mathbb{E}}_{(\vec{x},\vec{y})\sim\mathcal{D}(B)}[f(x_{t-1})\cdots f(x_0)\,v\otimes x_t\otimes y_t]\,.\end{aligned}$$

The second equality uses the inductive hypothesis, and the last two equalities use Eq. (4.4) for $\ell = t-1$ and $\ell = t$ respectively. □

Using Definition 4.3.6, we further define the operators for the distributions we wish to study.

Uniform Distribution Let us first capture, using this notation, the uniform distribution on walks on X starting from $x_0 \in V_x$. We define B_U where for each i, $B_i = A_X \otimes I_Y$ for every i. Then, for any ℓ, $(A_X \otimes I_Y)^\ell = A_X^\ell \otimes I_Y$. Therefore, we obtain that $\mathcal{D}_X(B_U)$ is the t-step random walk distribution on X i.e., $x_i \sim A_X^i x_0$.

The s-wide Distribution This is the distribution obtained by the s-wide walks as described in the earlier section. For $0 \le a \le b < s$, we define,

$$B[a,b] = \left(X_a\,\mathring{A}_Y, X_{a+1}\,\mathring{A}_Y, \cdots, X_b\,\mathring{A}_Y\right).$$

We can view this random walk as occurring in two steps. The first step picks an initial vertex $y_0 \in Y$ and the next picks the sequence of neighbors according to which we will

perform the walk on Y. To formalize this, let $A_Y = (1/d_2)\sum_{j=1}^{d_2} P_j$ where P_j are permutation matrices and let $J = (j_0, \cdots, j_{b-a}) \in [d_2]^{b-a+1}$. For a fixed sequence of permutations J, the conditional distribution (conditioned on J) is defined by,

$$B[a, b, J] = \left(X_a \overset{\circ}{P}_{j_0}, X_{a+1} \overset{\circ}{P}_{j_1}, \cdots, X_b \overset{\circ}{P}_{j_{b-a}}\right) .$$

We would like these two distributions, i.e., the uniform distribution and the s-wide distribution , to be the same and a graph Y is said to be *compatible* with respect to (X, ϕ), if for any fixed sequence, J, of a walk of length $\ell \leq s$, the distribution obtained on X via the uniform sampling of y_0, is the same as the usual ℓ-length walk on X from any fixed initial vertex, x_0. Thus, the randomness of sampling a vertex from Y is effectively *transferred* to a random walk on X.

Definition 4.3.8 (Compatible)**.** A graph Y is *compatible* with respect to (X, ϕ) if for every $0 \leq a \leq b < s$, $J \in [d_2]^{b-a+1}$ and $x_0 \in V_X$, we have[7],

$$\mathcal{D}_X(B[a, b, J], x_0) = \mathcal{D}_X(B_U, x_0) = A_X^{b-a+1} x_0 .$$

This *compatible* property is the same as the 0-pseudorandom property in [TS17]. We renamed it as it is more of a structural compatibility property than a pseudorandomness one. For the sake of completeness, we now prove that Cayley graphs are compatible with every locally invertible graph.

Lemma 4.3.9 ([TS17, Lemma 29])**.** *Let* $Y = \mathrm{Cay}(G^s, T)$ *where* $|G| = d_1$*. Then,* Y *is compatible with respect to any* X, ϕ*.*

Proof. Let $x_0 \in V_X$ be arbitrary and let $y_0 = (r_1, \cdots, r_s) \sim G^s$ be sampled uniformly. Since Y is a Cayley graph, the sequence of permutations, J, is equivalent to a sequence of generators $(t,^1, \cdots, t^s) \in T^s$, and the permutation is group multiplication, $y \mapsto t^k \cdot y$.

For each $1 \leq i \leq s$, let $t^i \cdot y_{i-1} = (r_1^i, \cdots, r_s^i)$. The walk evolves as follows,

7. It is important to note that $\mathcal{D}_Y(B[a, b, J]) \neq \mathcal{D}_Y(B_U)$.

$$x_i = X_i(x_{i-1}\,,\, t^i \cdot y_{i-1}) = x_{i-1}\,[r^i_i]$$
$$y_i = \phi(t^i \cdot y_{i-1}) = \left(r^i_1, \cdots, r^i_{i-1},\, \phi(r^i_i),\, r^i_{i+1}, \cdots, r^i_s\right).$$

Compatibility requires that $x_i \sim A^i_X x_0$, which is inductively implied if, for every i, r^i_i is uniform over G and independent of r^j_j for $j \neq i$.

This claim is true initially as y_0 is uniform over G^s. Assume it is true for y_{i-1}. Since t^i, ϕ are fixed permutations, they do not affect the uniformity of the distribution of r^i_k for any k. Since, ϕ acts only on the i^{th} component, the independence of r^i_k and r^i_i, guaranteed by inductive claim, implies the independence of r^i_k and $\phi(r^i_i)$. □

For a fixed x_0 and J, we say that y_0 *gives rise* to a sequence $\vec{z} = (z_1, \cdots, z_t)$ if $z_i = x_i$ where x_i is as defined in the above proof.

Observation 4.3.10. Fix $x_0 \in V_X$ and a sequence J. For any $\vec{z} = (z_1, \cdots, z_k)$, the number of y_0 that *gives rise* to $\vec{z}$ is d_2^{s-k}.

Proof. From the proof of Lemma 4.3.9, we see that enforcing $z_i = x_i$ for each i starting from $i = 1$, forces exactly r_i to be fixed. Thus, the remaining $(r_{k+1}, \cdots, r_s)$ are free. □

4.3.3 The s-wide Operator Norm Decay

We are now ready to establish the key technical lemma in the analysis of the s-wide replacement, which is an operator-valued generalization of the scalar version of present in [TS17].

Lemma 4.3.11 (Simulation Lemma (generalization of Lemma 26 from [TS17]))**.** *Let* $0 \leq s_1 \leq s_2 < s$. *For every pair of vectors* $z, z' \in \mathcal{X}_{\mathcal{H}}$, *we have,*

$$\left\langle \prod_{i=s_1}^{s_2} \left(\mathring{X}_i \mathring{A}_Y \mathring{\Pi}_f\right) \left(z \otimes \frac{1}{|V_Y|}\vec{1}\right), z' \otimes \vec{1} \right\rangle \;=\; \left\langle \left(\mathring{A}_X\, \Pi_f\right)^{s_2 - s_1 + 1} z, z' \right\rangle.$$

Proof. Let $z = \sum_x v_x \otimes x$ and $z' = \sum_x w_x \otimes x$. Since the expression is bilinear, it suffices to prove the equation for $v \otimes x_0$, $w \otimes x'$ for an arbitrary pair (x_0, x'). Let $t = s_2 - s_1 + 1$.

$$\prod_{i=s_1}^{s_2} \left(\mathring{X}_i \mathring{A}_Y \mathring{\Pi}_f \right) = \mathop{\mathbb{E}}_{(j_{s_1}, \cdots, j_{s_2}) \sim [d_2]^t} \left[\prod_{i=s_1}^{s_2} \left(\mathring{X}_i \mathring{P}_{j_i} \mathring{\Pi}_f \right) \right]$$

Therefore, we can fix $J = (j_{s_1}, \cdots, j_{s_2}) \in [d_2]^t$ and prove it for that. Recall the notation $\mathcal{W}_t$ that denotes the set of t-length walks on the graph X. Applying Lemma 4.3.7 to $B[s_1, s_2, J]$, we get,

$$\begin{aligned}
\prod_{i=s_1}^{s_2} \left(\mathring{X}_i \mathring{P}_{j_i} \mathring{\Pi}_f \right) \left(v \otimes x_0 \otimes \frac{1}{|V_Y|} \vec{1} \right) &= \mathop{\mathbb{E}}_{(\vec{x}, \vec{y}) \sim D(B[s_1, s_2, J])} [f(x_{t-1}) \cdots f(x_0)\, v \otimes x_t \otimes y_t] \\
&= \sum_{\vec{x} \in \mathcal{W}_t} \mathop{\mathbb{E}}_{y_0 \sim V_Y} [f(\vec{x})\, v \otimes x_t \otimes y_t] \mathbb{I}[y_0 \text{ gives rise to } \vec{x}], \\
&= \frac{d_1^{s-t}}{d_1^s} \sum_{\vec{x} \in \mathcal{W}_t} f(\vec{x})\, v \otimes x_t \otimes y_t,
\end{aligned}$$

where $f(\vec{x}) = f(x_{t-1}) \cdots f(x_0)$. The last equality uses Observation 4.3.10. Therefore, the conditioning on y does not change the distribution $\mathcal{D}_X$ and when we take inner products, we obtain,

$$\begin{aligned}
\left\langle \prod_{i=s_1}^{s_2} \left(\mathring{X}_i \mathring{P}_{j_i} \mathring{\Pi}_f \right) \left(v \otimes x_0 \otimes \frac{1}{|V_Y|} \vec{1} \right), w \otimes x' \otimes \vec{1} \right\rangle &= \frac{d_1^{s-t}}{d_1^s} \sum_{\vec{x} \in \mathcal{W}_t} \left\langle x_t, x' \right\rangle \langle f(x_{t-1}) \cdots f(x_0)\, v, w \rangle \\
&= \mathop{\mathbb{E}}_{\vec{x} \sim D_X(B[s_1, s_2, J])} \left[\left\langle x_t, x' \right\rangle \langle f(x_{t-1}) \cdots f(x_0)\, v, w \rangle \right].
\end{aligned}$$

We now use[8] Lemma 4.3.7 for B_U and take inner product to get,

$$\left\langle \left(\mathring{A}_X \Pi_f\right)^{s_2-s_1+1} (v \otimes x_0), w \otimes x' \right\rangle = \mathop{\mathbb{E}}_{\vec{x}\sim D_X(B_U)} \left[\langle x_t, x' \rangle \langle f(x_{t-1}) \cdots f(x_0) v, w \rangle \right].$$

From Lemma 4.3.9, we know that Y is compatible and thus, $\mathcal{D}_X(B[s_1, s_2, J]) = \mathcal{D}_X(B_U)$. Thus, the right-hand side (and so, the left-hand sides) of these two equations above are equal. □

The s-step Decay Just like the amplification in Section 4.2 was analyzed by studying the norm decay obtained in every two steps (*cf.*,Lemma 4.2.8), this amplification via the s-wide walks will be analyzed by bounding the norm decay for steps of length s using Lemma 4.3.11 similarly to [BATS08, TS17]. We will use the shorthand $L_i := \mathring{X}_i \mathring{\Pi}_f \mathring{A}_Y$.

The goal is to bound $\|L_{s-1} \cdots L_0\|_{op}$ which controls the bias of the set obtained by s-long, s-wide walks (*cf.*,proof of Eq. (4.2)). Equivalently, we will bound $\langle (\prod_i L_i) v_0, w_s \rangle$ for any unit vectors[9] $v_0, w_s \in \mathcal{XY}_{\mathcal{H}}$. We will use the orthogonal decomposition,

$$\mathcal{XY}_{\mathcal{H}} := \mathcal{X}_{\mathcal{H}} \otimes \mathbb{C}[V_Y] = \mathcal{XY}^{\|}_{\mathcal{H}} \oplus \mathcal{XY}^{\perp}_{\mathcal{H}} \text{ where } \mathcal{XY}^{\|}_{\mathcal{H}} := \mathrm{span}\left\{ z \otimes \vec{1} \mid z \in \mathcal{X}_{\mathcal{H}} \right\}.$$

For $i \geq 1$, we inductively define the vectors v_i, w_i, z_i and bound their norms[10],

$$v_i = L_{i-1} v^{\perp}_{i-1}, \qquad z_{s-i} = \left(\mathring{X}_{s-i} \mathring{\Pi}_f\right)^* w_{s-i+1}, \qquad w_{s-i} = \left(\mathring{A}_Y\right)^* z^{\perp}_{s-i} \tag{4.5}$$

$$\|v_i\| \leq \lambda(Y)^i, \qquad \|z_{s-i}\| \leq \lambda(Y)^{i-1}, \qquad \|w_{s-i}\| \leq \lambda(Y)^i. \tag{4.6}$$

Lemma 4.3.12. *For any v_0, w_s and $0 \leq r \leq s-2$ we have,*

8. As we only want to work with the space $\mathcal{X}_{\mathcal{H}}$ here, we can assume in the application of the lemma that $|V_Y| = 1$. Else, one could directly apply Eq. (4.1) and use the observation that $\mathcal{D}_X(B_U)$ is the same as the random walk distribution on X.

9. Here we deviate from our notation and use v, w for vectors in $\mathcal{XY}_{\mathcal{H}}$.

10. By definition $\|v_i\| \leq \|\mathring{A}_Y v^{\perp}_{i-1}\| \leq \lambda(Y)\|v_{i-1}\|$.The computation is similar for w and z.

$$L_{s-1}\cdots L_0 v_0 \;=\; v_s \;+\; \sum_{i=0}^{s-1} L_{s-1}\cdots L_i v_i^{\|}$$

$$L_{s-1}^* w_s = w_{s-1} + z_{s-1}^{\|}$$

$$L_r^*\cdots L_{s-1}^* w_s \;=\; w_r \;+\; z_r^{\|} \;+\; \sum_{i=r+1}^{s-1} L_r^*\cdots L_{i-1}^* z_i^{\|}$$

Proof. The lemma follows readily from a calculation, and we omit its proof. □

Theorem 4.3.13 (Operator Generalization of Theorem 24 [TS17])**.** *Let* X *be any* d_1*-regular graph and* Y *be a Cayley graph on* $\mathbb{F}_2^{s\log d_1}$*. Let* $s\mathcal{W}_t$ *be the collection of* t*-length* s*-wide walks on the s-wide replacement product on* X *and* Y*. For any operator-valued function* f *on* V_X*, such that* $\max_{x\in V_X}\|f(x)\|_{op} \le 1$ *and* $\left\|\mathbb{E}_{x\in V_X}[f(x)]\right\|_{op} := \lambda_0 \le \lambda(Y)^2 - 2\lambda(X)$*,*

$$\left\|\mathop{\mathbb{E}}_{\vec{x}\in s\mathcal{W}_t}[f(x_t)\cdots f(x_0)]\right\|_{op} \;\le\; \left(\lambda(Y)^s + s\cdot\lambda(Y)^{s-1} + s^2\cdot\lambda(Y)^{s-3}\right)^{\lfloor t/s\rfloor}.$$

Proof. Using Lemma 4.3.7, we can repeat the proof of Eq. (4.2) to see that,

$$\left\|\mathop{\mathbb{E}}_{\vec{x}\in s\mathcal{W}_t}[f(x_t)\cdots f(x_0)]\right\|_{op} \;\le\; \|L_t\cdots L_0\|_{op} \;\le\; \|L_{s-1}\cdots L_0\|_{op}^{\lfloor t/s\rfloor}.$$

We now use Lemma 4.3.12 to bound this operator norm.

$$\begin{aligned}
\langle L_{s-1}\cdots L_0 v_0, w_s\rangle \;&=\; \langle v_s, w_0\rangle \;+\; \sum_{r=0}^{s-1}\left\langle L_{s-1}\cdots L_r v_r^{\|}, w_s\right\rangle \\
&=\; \langle v_s, w_s\rangle \;+\; \sum_{r=0}^{s-1}\left\langle v_r^{\|}, L_r^*\cdots L_{s-1}^* w_s\right\rangle \\
&=\; \langle v_s, w_s\rangle \;+\; \sum_{i=0}^{s-1}\left\langle v_r^{\|}, w_r + z_r^{\|}\right\rangle \;+\; \sum_{r=0}^{s-2}\sum_{i=r+1}^{s-1}\left\langle v_r^{\|}, L_r^*\cdots L_{i-1}^* z_i^{\|}\right\rangle \\
&=\; \langle v_s, w_s\rangle \;+\; \sum_{i=0}^{s-1}\left\langle v_r^{\|}, z_r^{\|}\right\rangle \;+\; \sum_{r=0}^{s-2}\sum_{i=r+1}^{s-1}\left\langle v_r^{\|}, L_r^*\cdots L_{i-1}^* z_i^{\|}\right\rangle.
\end{aligned}$$

The last step uses $\left\langle v_r^{\|}, w_r\right\rangle = \left\langle \mathring{A}_Y\, v_r^{\|}, z_r^{\perp}\right\rangle = 0$. Using Eq. (4.6), we get $\left\langle v_r^{\|}, z_r^{\|}\right\rangle \le \lambda(Y)^{s-1}$.

To bound the last term, we finally use Lemma 4.3.11. Let $v_r^{\|} = v'_r \otimes \vec{1}$, and $z_i^{\|} = z'_i \otimes \frac{1}{|V_Y|}\vec{1}$. Then,

$$\begin{aligned}
\left\langle v_r^{\|}, L_r^* \cdots L_{i-1}^* z_i^{\|} \right\rangle &= \left\langle v'_r, \left(\mathring{A}_X \Pi_f\right)^{i-r} z'_i \right\rangle && \text{(Using Lemma 4.3.11)} \\
&\leq \left\| \left(\mathring{A}_X \Pi_f\right)^{i-r} \right\|_{op} \|z'_i\| \|v'_r\| \\
&\leq \lambda(Y)^{2\lfloor \frac{i-r}{2} \rfloor} \lambda(Y)^{r+s-i-1} \leq \lambda(Y)^{s-3},
\end{aligned}$$

where the penultimate inequality uses Theorem 4.2.1 and plugs in the assumption that $2\lambda(X) + \left\| \mathbb{E}_{x \in V_X}[f(x)] \right\|_{op} \leq \lambda(Y)^2$. Substituting this back in our expression above gives us the result. □

This concludes the proof of the key technical theorem, i.e., the operator amplification as outlined in Theorem 4.1.1. One needs to initialize this product and verify that $|W|$ is indeed as small as we need, $|W| \leq O_{|S|}(\lambda^{-2-o(1)})$. However, this instantiation closely follows Ta-Shma's [TS17], and we relegate the computations to Appendix A.1. Apart from this calculation, this finishes the proof of Theorem 1.2.2

Part II

Applications

CHAPTER 5

PSEUDORANDOM OBJECTS

Laqa: Seeing these mulberry trees blossom reminds me of home.

Janki: Me too, I am glad TTIC has these trees around. I have been meaning to ask, you often speak about expanders with symmetry. Why do we need this added symmetry?

Laqa: Okay, so a classical code is a subspace that you can construct out of a graph.

Janki: Yes, and if the graph is an expander, you get a code with a large distance.

Laqa: Alright, but a quantum CSS code is slightly more complicated. To construct it, we need a "2-dimensional chain complex" whereas a graph is a 1-D complex. One way people constructed these "2-D" complexes was via topology, but such constructions often do not have expansion, and the distance is not great.

Janki: So instead, we want a bottom-up combinatorial method that starts from two expanders?

Laqa: Exactly! One can get a 2-D complex by tensoring two 1-D complexes, but such a tensor product based code can never have a distance better than $\sqrt{n}$.

Janki: Just like the cartesian product of two complete graphs will have eigenvalue $\sqrt{n}$?

Laqa: Yes, much like that. An exciting recent development has been the idea of getting around this by quotienting this tensor product. The hope is that the complex shrinks more than the distance. But to do that, we need the base graphs to have symmetry. In general, symmetry helps when we need more structure in pseudorandom objects like a *quantum expander*, or a *dimension expander*.

5.1 Explicit Quantum and Classical Codes

The area of quantum error correction has had tremendous progress in recent years, particularly in the construction of quantum low-density parity-check (QLDPC) codes. These are codes where membership can be tested via checks acting only on a small number of qubits, which is a very useful property for physical implementations. Since the classic toric code by Kitaev [Kit97], there has been a flurry of activity [TZ14, EKZ20, KT21, HHO21, PK21, PK22, LZ22, DHLV23] culminating in the breakthrough construction of *asymptotically good* QLDPC codes by Panteleev and Kalachev [PK22]. Specifically, the above constructions yield a special form of quantum code known as Calderbank–Shor–Steane (CSS) code, first described in [CS96, Ste96]. These can be specified by a pair of subspaces $\mathcal{C}_X, \mathcal{C}_Z \subseteq \mathbb{F}_2^n$ satisfying $\mathcal{C}_Z^\perp \subseteq \mathcal{C}_X$ (which also implies $\mathcal{C}_X^\perp \subseteq \mathcal{C}_Z$). The quantum CSS codes are LDPC when $\mathcal{C}_X^\perp$ and $\mathcal{C}_Z^\perp$ have generating sets consisting of sparse vectors. The code $\mathcal{C} = (\mathcal{C}_X, \mathcal{C}_Z)$ is said to have blocklength n, with distance d and the dimension k defined as,

$$d = \min\left\{|c| \mid c \in (\mathcal{C}_X \setminus \mathcal{C}_Z^\perp) \cup (\mathcal{C}_Z \setminus \mathcal{C}_X^\perp)\right\} \quad \text{and} \quad k = \left(\dim(\mathcal{C}_X) - \dim(\mathcal{C}_Z^\perp)\right).$$

A code is called *asymptotically good* if $d, k = \Omega(n)$. In all these works, the code is constructed using a pair of expander graphs, each with symmetries of some group G. This construction connects symmetric expanders with the area of quantum error correction and provides a very concrete motive to study the explicit construction of such graphs.

Quasi-Cyclic Codes Graphs with symmetry also yield classical linear codes with symmetry. While such a symmetry does not improve its distance or dimension, it is often helpful for algorithmic purposes. For example, the cyclic symmetry of *cyclic codes*, i.e., codes that are invariant under the action of $\mathbb{Z}_N$ where N is the block length, leads to efficient encoding and decoding algorithms. Babai, Shpilka, and Stefankovich [BSS05] showed that cyclic codes cannot be *asymptotically good LDPC codes*.

Quasi-cyclic codes are a generalization of cyclic codes in which symmetry is only under a subgroup, $\mathbb{Z}_\ell \leqslant \mathbb{Z}_N$, where $N = n\ell$. The parameter ℓ is called the *circulant size*, and the closer it is to N, the closer the code is to being cyclic. Interestingly, with this relaxation, good quasi-cyclic codes are known to exist, as shown by Chen, Peterson, and Weldon [CPW69]. More recently, Bazzi and Mitter [BM06] gave a randomized construction for any constant $n > 2$ and showed that it attains Gilbert–Varshamov bound with rate $1/n$. Quasi-cyclic codes have been extensively studied and are very useful in practice (e.g., their LDPC counterparts are part of the 5G standard of mobile communication [LBM+18]).

The Pantaleev-Kalachev construction The work [PK21] gives a construction of good quasi-cyclic codes and quantum CSS codes. The follow-up work [PK22], when applied to the Abelian group $\mathbb{Z}_\ell$, improves the dimension of the quantum CSS code.

Theorem 5.1.1 ([PK21, PK22])**.** *Let X be $\mathbb{Z}_\ell$-lift of a d-regular graph on n-vertices with $\lambda_2(X) \leq \varepsilon \cdot d$. If $\varepsilon > 0$ is sufficiently small and d is a sufficiently large constant, then,*

- *There exists a good quasi-cyclic LDPC code of blocklength $\Theta(n\ell)$ and circulant size $\Theta(\ell)$.*
- *There exists an LDPC quantum CSS code of distance $\Theta_{\varepsilon,d}(\ell)$ and dimension $\Theta(n\ell)$.*

For this application, the constant degree regime is essential for two reasons. The locality of the code is essentially d, and thus, it has to be constant for it to be LDPC. Moreover, this construction relies on a brute-force search over subspaces of $\mathbb{F}_2^d$.

To achieve these, [PK21] picks a d-regular expander on n vertices and creates a random ℓ-lift which is expanding with high probability Theorem 2.1.3.

The distance achieves the almost-linear bound only when the lift is large, and thus, lifts of exponential size are preferred. By the upper bound in Theorem 2.1.3, better than exponential size lifts break expansion for abelian groups.

5.1.1 Derandomized Codes

We restate our main results about the explicit construction of graphs with abelian lifts. This summarizes Theorem 3.0.1 and Theorem 2.2.1.

Theorem 1.2.1 (Explicit Abelian Lifts)**.** *For large enough* n *and constant degree* $d \geq 3$*, given an abelian group* G*, and any fixed constant* $\varepsilon \in (0,1)$*, we can construct a* d*-regular graph* X *on* $\Theta(n|G|)$ *vertices, in deterministic polynomial time, such that,*

1. X *is* G*-lift of a graph* X_0 *on* $\Theta(n)$ *vertices. Thus,* $G \subseteq \mathrm{Aut}(X)$*.*
2. *If* $|G| \leq \exp\left(n^{\delta(d,\varepsilon)}\right)$*, then* $\lambda_u(X) \leq 2\sqrt{d-1} + \varepsilon$*.*
3. *If* $|G| \leq \exp\left(n^{\delta}\right)$ *and also* $d \geq d_0(\varepsilon)$*, then* $\lambda_u(X) \leq \varepsilon \cdot d$*.*
4. *If* $|G| \leq \exp\left(cnd^{-\frac{1}{2}}\right)$*, then* $\lambda_u(X) \leq O(\sqrt{d}\log d)$*.*
5. *If* $|G| = \exp\left(cnd^{\delta}\right)$ *for* $\delta \in [-1/2, 1)$*, then* $\lambda_u(X) \leq O\left(d^{\frac{2+\delta}{3}}\log d\right)$*.*

While the corollary follows straightforwardly from Theorems 1.2.1 and 5.1.1, we show the computations for completeness.

Corollary 5.1.2. *We have explicit polynomial time construction of each of the following,*

1. *Good quasi-cyclic LDPC code of block length* N *and any circulant size up to* $N/\mathrm{polylog}(N)$ *or* $\Theta(N/\log(N))$*.*
2. *Quantum LDPC code with distance* $\Omega(N/\log(N))$ *and dimension* $\Omega(N)$*.*

Proof. We use Theorem 1.2.1 in the exponential regime to construct X explicitly. When $\ell = \exp(\Theta(n))$, $N = n\ell$ and we get quantum LDPC codes with distance $\Theta(\ell) = \Theta(N/\log N)$.

For quasi-cyclic codes, we can set $\ell \leq 2^{n^{\delta_0}}$ with some fixed $\delta_0 \in (0,1)$, and explicitly construct X which is a $\mathbb{Z}_\ell$-lift by Theorem 1.2.1. By Theorem 5.1.1, there is a code with circulant size $\Theta(\ell)$ and $\log(N) \leq \log n + n^{\delta_0} \leq 2n^{\delta_0}$ (for n sufficiently large). Thus, the construction works for circulant sizes $\ell = O\left(N/(\log N)^{1/\delta_0}\right)$. □

5.2 Other Pseudorandom Objects

We will now discuss some applications of the operator amplification technique from Chapter 4, which allows us to improve other pseudorandom objects. All the "pseudorandom" objects below are expanders (with various structural properties), and for each of these, we amplify their spectral bound to almost Ramanujan. We stress that our amplification preserves the underlying structure and, therefore, produces another object with the same properties. We first give an overview of the results, and the details will follow.

5.2.1 Overview of Results

General Expanders A very interesting and useful application is that the amplification for Cayley graphs implies an expansion result for general families of (regular) expander graphs. This highlights how studying structured graphs (like Cayley graphs) can shed light on general graphs.

Theorem 5.2.1 (Amplifying General Expanders)**.** *Let* $\{X_i\}_{i\in\mathbb{N}}$ *be a family of* (d_0, λ_0)*-expanders where* $\lambda_0 < 1$ *is a constant. For any (target)* $\lambda \in (0,1)$ *and* X_i*, we can explicitly*[1] *construct a* (d, λ)*-expander,* X'_i*, on the same vertex set, where* $d = O(d_0/\lambda^{2+o_\lambda(1)})$*. Moreover, the construction is local in the sense that edges in* X'_i *correspond to short walks in* X_i*.*

Quantum Expanders Roughly speaking, a quantum expander is an operator defined by d complex matrices, whose (linear) action on quantum states has a constant spectral gap. Quantum expanders were defined in [AS04, BASTS08, Has07a], and Hastings [Has07b] showed that the Ramanujan bound also applies to them. Existing explicit constructions are far from the Ramanujan bound. In [Har07], Harrow gave a generic construction using expanding Cayley graphs, which is explicit if the group has a large irreducible representation and admits efficient Quantum Fourier Transform (QFT). Both these conditions are

1. See Definition 1.3.6

satisfied by the symmetric group Sym_n using the generating family by Kassabov [Kas07] and the QFT algorithm by Beals [Bea97].

By amplifying the expansion of the generators of [Kas07], we give the first explicit family of almost Ramanujan quantum expanders.

Corollary 5.2.2 (Explicit Almost Ramanujan Quantum Expanders)**.** *For every* $\lambda \in (0,1)$*, there is an explicit infinite family of (efficient)* $(O(1/\lambda^{2+o(1)}), \lambda)$*-quantum expanders.*

Monotone Expanders Monotone expanders are expanders whose edge set can be decomposed into a constant number of *monotone* partial maps on $[n]$. Bourgain and Yehudayoff [BY13] gave the only known explicit construction of monotone expanders with *constant* degree. There are two natural notions of degree for a monotone expander. The usual vertex degree and the number of monotone maps. Our almost Ramanujan trade-off is with respect to the vertex degree (and the monotone degree is polynomial in the vertex degree).

Corollary 5.2.3 (Almost Ramanujan Monotone Expanders)**.** *For every* $\lambda > 0$*, there is an explicit family* $\{X_i\}_{i\in\mathbb{N}}$ *of (vertex)* d*-regular* $d^{O(1)}$*-monotone expanders with* $d = O(1/\lambda^{2+o(1)})$ *and* $\lambda(X_i) \leq \lambda$*.*

The approach is similar to that used for Theorem 5.2.1; we express it as a sum of permutation matrices and amplify their expansion, obtaining the following result. It would be really interesting to obtain an almost Ramanujan trade-off with respect to the monotone degree.

Dimension Expanders Loosely speaking, dimension expanders (over any field $\mathbb{F}$) are a linear algebraic extension of expanders: a collection of d linear maps on $\mathbb{F}^n$, which significantly *expands* (the span of) any vector space of dimension below $n/2$. They were defined by Barak et al. in [BISW01]. Over complex numbers, any quantum expander is a

dimension expander. More generally, Dvir and Shpilka [DS09] proved that a monotone expander directly yields a dimension expander over every field. We give spectral almost Ramanujan expanders with the additional property of dimension expansion. Additionally, if the starting dimension is small enough, then we achieve almost doubling of the starting dimension See Corollary 5.2.25 for a precise statement.

Improving the Kazhdan Constant The *Kazhdan constant* $\mathcal{K}(G, S)$ of a finitely generated group G, with respect to a generating set S, is a quantitative version of Property (T) which has been used to construct explicit expanders (e.g., Margulis [Mar88]). We show that this can be amplified by considering a slightly different version called the *average Kazhdan constant* which directly relates to the bias of the set S. This is interesting as typically the bound on the Kazhdan constant is used to construct expanders, but here, we construct expanding generating sets to improve the constant!

Corollary 5.2.4 (Amplifying Average Kazhdan Constant)**.** *Let* G *be a finitely generated group and* S *a finite set of generators such that the average Kazhdan constant* $\overline{\mathcal{K}}(G, S)$ *is equal to* $2 \cdot (1 - \lambda_0)$ *for some constant* $\lambda_0 \in (0, 1)$*. For every* $\lambda \in (0, 1)$*, there is a set* $S' \subseteq G$ *such that*

1. $\overline{\mathcal{K}}(G, S') \geq 2 \cdot (1 - \lambda)$*, and thus,* $\mathcal{K}(G, S') \geq 2 \cdot (1 - \lambda)$*.*
2. $|S'| = O_{\lambda_0}(|S|/\lambda^{2+o(1)})$*, and*
3. S' *can be found in time* $\mathrm{poly}(|S|/\lambda)$ *assuming an oracle for group operations on* G*.*

The improved constants and the generating sets have algorithmic implications, and we mention two of them.

- *Dimension expanders* - Lubotzky and Zelmanov [LZ08] showed that the image of a generating set of a group under an irreducible representation gives a dimension expander and its expansion is controlled by its Kazhdan constant.

- *Product replacement algorithm* - uses random walks on k-tuples of groups elements. Lubotzky and Pak [LP00] showed that the mixing time of the algorithm relates to the Kazhdan constant of certain lattice groups like $\mathrm{SL}_n(\mathbb{Z})$, assuming Property (T). This crucial assumption was proven in complete generality[2] recently by Kaluba, Kielak and Nowak [KKN21]. In particular, we have a mixing time bound of $\frac{4\log|G|}{\mathcal{K}(G,S)^2}$.

We can improve both results by using our amplified generating set (Corollary 5.2.19).

5.2.2 *Permutation Amplification*

The *defining representation* - $(\rho_{\mathrm{def}}(\sigma), \mathbb{C}^n)$ for Sym_n is defined as the representation that maps a permutation to the matrix defining it. More formally, $\rho_{\mathrm{def}}(\sigma)e_i = e_{\sigma(i)}$ for every unit basis vector e_i of $\mathbb{C}^n$. It is a fact that $\mathcal{V}_{\mathrm{def}} = \mathcal{V}_{\mathrm{triv}} \oplus \mathcal{V}_{\mathrm{standard}}$ where $\mathcal{V}_{\mathrm{standard}}$ is an irreducible non-trivial representation. Note that if we are given a set $\{P_1, \cdots, P_r\}$ of permutation matrices acting on $\mathbb{C}^n$, we can identify a set $S = \{\sigma_1, \cdots, \sigma_r\} \subseteq \mathrm{Sym}_n$ such that $\rho_{\mathrm{def}}(\sigma_i) = P_i$.

Corollary 5.2.5 (Permutation Amplification)**.** *Let* $P = \{P_1, \cdots, P_r\}$ *be a collection of permutation matrices such that* $\lambda(\mathbb{E}_{i\sim[r]}[P_i]) \leq \lambda_0$*. Then, for any* $\lambda \in (0,1)$*, we can explicitly construct a collection* P' *such that*

1. $\lambda(\mathbb{E}_{M\sim P'}[M]) \leq \lambda$,
2. $|P'| \leq O\left(|P|/\lambda^{2+o(1)}\right)$ *and*
3. *each* $P'_i \in P'$ *is a product of at most* $O_{\lambda_0}(\log(1/\lambda))$ *many matrices from* P.

Proof. Let $P_i = \sigma_i$. Applying Theorem 4.3.13 to the set $S = \{\sigma_i\}$ we get a larger set of permutations, S' of the form $\sigma' = \sigma_{i_1} \circ \cdots \circ \sigma_{i_k}$ where $k = O_{\lambda_0}(\log(1/\lambda))$. By the

2. [KKN21] prove that $\mathrm{Aut}(F_n)$, the automorphism group of the free group generated by n elements, has Property (T). This implies Property (T) for quotients of $\mathrm{Aut}(F_n)$, which includes $\mathrm{SL}_n(\mathbb{Z})$.

decomposition of the defining representation, we have that

$$\begin{aligned}\mathrm{Spec}\Big(\mathop{\mathbb{E}}_{M\sim P'}[M]\Big) &= \mathrm{Spec}\Big(\mathop{\mathbb{E}}_{\sigma'\sim S'}\big[\rho_{\mathrm{def}}(\sigma')\big]\Big)\\ &= \{1\}\cup\mathrm{Spec}\Big(\mathop{\mathbb{E}}_{\sigma'\sim S'}\big[\rho_{\mathrm{standard}}(\sigma')\big]\Big)\,.\end{aligned}$$

where the 1 corresponds to the eigenvalue from the trivial representation. Since the operator amplification reduces the bias of every non-trivial irreducible representation, it also does so for $\mathcal{V}_{\mathrm{standard}}$. □

5.2.3 Arbitrary Expanders via Permutation Amplification

We can make any family of bounded degree expander graphs into an almost Ramanujan family while preserving their adjacency structure. First, we recall König's theorem that says that the adjacency matrix of a d-regular graph can be expressed in terms of permutation matrices.

Theorem 5.2.6 (König)**.** *Let A_X be the normalized adjacency matrix of a d-regular n-vertex simple graph X. Then, there exists d permutation matrices $P_1,\ldots,P_d\in\mathbb{R}^{n\times n}$ such that*

$$A_X=\frac{1}{d}\sum_{j=1}^{d}P_j\ .$$

Claim 5.2.7. *The permutations in Theorem 5.2.6 can be found in time $\mathrm{poly}(n)$.*

Proof. We view A_X as encoding the adjacency relation of a bipartite graph with vertex bipartition $(A=V(X),B=V(X))$. This bipartite graph is d-regular so it has at least one perfect matching M, which can be found in $\mathrm{poly}(n)$ time. We remove this matching M, obtaining a $(d-1)$-regular graph, and repeat until the resulting graph is empty. □

Our general transformation into an almost Ramanujan bound follows by using Claim 5.2.7 to obtain an initial set of permutation matrices and amplify then using Corollary 5.2.5.

Theorem 5.2.8 (Main I (Formal version of Theorem 5.2.1))**.** *Let* $\{X_i\}_{i\in\mathbb{N}}$ *be a family of* d_0*-regular* λ_0*-expanders with constant* $\lambda_0 < 1$*. For any* $\lambda \in (0,1)$ *and any expander* X_i*, we can deterministically compute a* d*-regular* λ*-expander* X'_i *with* $d = O_{\lambda_0}(d_0/\lambda^{2+o(1)})$ *in time* $\mathrm{poly}(|V(X_i)|)$*. Moreover, the construction is local in the sense that edges in* X'_i *correspond to short walks in* X_i*. More precisely, if the adjacency matrix of* X_i *is*

$$A_{X_i} = \frac{1}{d_0}\sum_{j=1}^{d_0} P_j,$$

where $P_1, \ldots, P_{d_0}$ *are permutation matrices, then the adjacency matrix of* X'_i *is*

$$A_{X'_i} = \frac{1}{d}\sum_{j=1}^{d} P'_j,$$

where each P'_j *is the product of at most* $k = O_{\lambda_0}(\log(1/\lambda))$ *permutation matrices among* $P_1, \ldots, P_{d_0}$*.*

5.2.4 Explicit Almost Ramanujan Quantum Expanders

Quantum expanders were defined in [AS04, BASTS08, Has07a] and have found many applications in quantum information theory. For instance, they can be used in the construction of designs and gates sets [HH09], in quantum statistical zero-knowledge (QSZK) [BASTS08], in detecting EPR pairs [AHL+14] and in the study of *entanglement* [Has07a].

While a usual degree-d expander graph $X = (V, E)$ is given by d permutation matrices acting on a vector space $\mathbb{C}[V]$, a quantum expander is given by d (suitable) linear operators acting on quantum states (i.e., PSD matrices of trace 1). The normalized adjacency matrix of a λ-expander shrinks the ℓ_2-norm of vectors orthogonal the all ones function by a factor of λ. Similarly, a quantum expander shrinks the Frobenius norm of PSD matrices orthogonal [3] to the identity matrix (the quantum analog of the all-ones function) by a factor of λ (the quantum expansion parameter).

3. With respect to the Hilbert–Schmidt inner product.

Definition 5.2.9 (Quantum Expander [AHL+14])**.** The (super) operator $\Phi : \mathbb{C}^{N\times N} \to \mathbb{C}^{N\times N}$ is an (N, d, λ) quantum expander if

· "Degree" – The operator Φ can be expressed as a sum of d linear operators as follows, $\Phi(\rho) = \sum_{i=1}^{d} B_i \rho B_i^\dagger$ where[4] $\sum_{i=1}^{d} B_i^\dagger B_i = I_N$.

· "Expansion" – The second largest eigenvalue[5] of Φ as a linear map is $\leq \lambda$.

In [Has07b], Hastings showed that the Ramanujan bound also applies to quantum expanders, and that d random unitaries get arbitrarily close to the bound. However, such a construction cannot be efficiently implemented and thus used in applications like [AHL+14], which rely on existing explicit constructions (e.g., [BASTS08, Har07]) that are far from the Ramanujan bound and thus give sub-optimal results.

Harrow [Har07] proved that one can construct a quantum expander using an expander Cayley graph over a group for which efficient Quantum Fourier Transform (QFT) is known.

Theorem 5.2.10 (Harrow [Har07])**.** *Let G be a group and $S \subseteq G$ be a multiset such that $\mathrm{Cay}(G, S)$ is a λ-spectral expander. Let V^μ be an irreducible representation of G of dimension N. Then, there exists an $(N, |S|, \lambda)$-quantum expander. Furthermore, if the group G admits an efficient QFT and $\log N = \Omega(\log|G|)$, then the quantum expander is explicit.*

Until now, we did not have almost-Ramanujan expanders over such a group. Since the symmetric group admits such a QFT algorithm [Bea97], we deduce the existence of explicit families of almost Ramanujan quantum expanders by applying our amplification to the Cayley graphs over the symmetric group due to Kassabov [Kas07].

Corollary 5.2.11 (Explicit Almost Ramanujan Quantum Expanders)**.** *For every $\lambda \in (0, 1)$, there is an explicit infinite family of (efficient) $(O(1/\lambda^{2+o(1)}), \lambda)$-quantum expanders.*

4. A useful special case is when each B_i is a (normalized) unitary.

5. If ρ satisfies $\mathrm{Tr}(\rho) = 0$, then $\|\Phi(\rho)\|_2 \leq \lambda\|\rho\|_2$, where $\|\rho\|_2 := \sqrt{\mathrm{Tr}(\rho^\dagger\rho)}$.

5.2.5 *Explicit Almost Ramanujan Monotone Expander*

We now show how to obtain almost Ramanujan monotone expanders starting from the explicit construction in Bourgain and Yehudayoff [BY13]. First, we recall the definition of a monotone graph. All graphs we consider are undirected.

Definition 5.2.12 (Monotone partial map)**.** A partial map $f : [n] \to [n]$ is monotone if for every pair $\{x, y\}$ for which f is defined, if $x < y$, we have $f(x) < f(y)$.

Definition 5.2.13 (Monotone Graph)**.** A bipartite graph $X = ([n]_A \sqcup [n]_B, E)$ is a d-monotone graph if there are d monotone partial maps $f_1, \ldots, f_d : [n] \to [n]$, such that the edges set E is the following disjoint union,

$$E = \bigsqcup_{i=1}^{d} \{(v_A, f_i(v)_B) \mid v \in \text{Domain}(f_i)\}.$$

We observe that there are two notions of degree of a monotone graph: the usual vertex degree and the number of monotone functions. Clearly, if a graph is d-monotone, all vertex degrees are at most d. The converse is not necessarily true; for example, the complete bipartite graph on 2 vertices on each side, $K_{2,2}$, has vertex degree 2, but the graph is not 2-monotone. We stress that our almost Ramanujan bound is with respect to the usual notion of vertex degree (and keeps the number of monotone maps polynomial in the vertex degree).

Definition 5.2.14 (Monotone Vertex Expander)**.** We say that $X = (A = [n]_A \sqcup B = [n]_B, E)$ is a d-monotone expander if it is a d-monotone graph and there exists $\delta > 0$ such that for all $A' \subseteq A$ with $|A| \leq n/2$, we have $|\partial(A')| \geq (1 + \delta)|A'|$, where $\partial(A')$ is the set of vertices in B adjacent to A'.

Theorem 5.2.15 (Bourgain and Yehudayoff [BY13])**.** *There is an explicit family* $\{X_n\}_{n \in \mathbb{N}}$ *of* d*-monotone vertex expanders with* $d = \Theta(1)$.

We will work with a spectral definition of a monotone expander. For a bipartite graph X, we define its *biadjacency matrix*, B_X such that the adjacency matrix $A_X = \begin{pmatrix} 0 & B_X \\ B_X^T & 0 \end{pmatrix}$. Precisely, for a monotone graph $X = ([n]_A \sqcup [n]_B, E)$, we have $(B_X)_{ij} = \mathbb{I}[(i_A, j_B) \in E]$. Note that if X is d-regular, then B_X is d-regular. We will define the graph X via B_X throughout.

Definition 5.2.16 (Spectral Monotone Expander)**.** We say that a d-monotone graph, X, is a λ-spectral monotone expander if $\lambda(X) = \max\{|\lambda_2(B_X)|, |\lambda_n(B_X)|\} < \lambda$.

It is well-known that starting from a monotone expander (not necessarily a vertex regular graph), we can add monotone partial functions to obtain a monotone graph of regular (vertex) degree that is still expanding. We use this to establish the following,

Corollary 5.2.17. *There is explicit family $\{X_n\}_{n\in\mathbb{N}}$ of d_0-regular $2d_0$-monotone expanders with $\lambda(X_n) \le \lambda_0 < 1$ and $d_0 = \Theta(1)$. Furthermore, the unnormalized adjacency matrix of X_n can be written as a sum of d_0 permutation matrices, each corresponding to two monotone maps.*

Proof. Let $\{X'_n\}_{n\in\mathbb{N}}$ be the family in Theorem 5.2.15. Let $X = X'_n$ be a fixed d_0-regfular graph that is also d_0-monotone expander with the maps $\{f_i\}$.

For each monotone function f_i, we define its "complement", $\overline{f}_i$, as the (unique) monotone partial function $\overline{f}_i$ such that $f_i \cup \overline{f}_i$ is a total function. Let Y be the $2d_0$-monotone graph corresponding to the maps $\{f_i, \overline{f_i}\}$. Then, we have

$$B_Y = \sum_{i=1}^{d_0} P_i \,,$$

where $P_i = M_{f_i} + M_{\overline{f_i}}$ and $(M_{f_i})_{x,y} = [f_i(x) = y]$.

Each matrix P_i is a permutation matrix as $f_i \cup \overline{f}_i$ is a total function. Adding more maps preserves the constant vertex expansion parameter, which (together with having constant

vertex degree) implies constant spectral expansion bounded away from 1 (see [Vad12, Theorem 4.9]). Thus, $\{Y_n\}_{n\in\mathbb{N}}$ is the required family. □

In the amplification process, we will be multiplying permutation matrices rather than just composing monotone maps since the latter operation can result in a map with an empty domain. We now establish the derandomized spectral amplification of monotone expanders.

Corollary 5.2.18 (Almost Ramanujan Monotone Expanders)**.** *For every $\lambda > 0$, there is an explicit family $\{X_i\}_{i\in\mathbb{N}}$ of (vertex) d-regular $d^{O(1)}$-monotone expanders with $d = O(1/\lambda^{2+o(1)})$ and $\lambda(X_i) \le \lambda$.*

Proof. Let $\{X'_n\}_{n\in\mathbb{N}}$ be the family in Corollary 5.2.17. Fix $X = X'_n$ and let $P_1, \dots, P_{d_0} \in \mathbb{R}^{n\times n}$ be the permutation matrices guaranteed by Corollary 5.2.17, where each $P_i = M_{f_i} + M_{\overline{f_i}}$. Use Corollary 5.2.5 to obtain a collection P' of size $|P'| = d := O(1/\lambda^{2+\beta})$ such that,

$$P' = \left\{\sigma \mid \sigma = P_{i_1} \cdots P_{i_k} \text{ for some } i_i, \cdots, i_k \in [d_0]\right\}.$$

Our final bipartite monotone graph will be Y given by $B_Y = \sum_{\sigma\in P'} \sigma$. To compute its monotone degree, we observe that,

$$\begin{aligned} P_{i_1} \cdots P_{i_k} &= \sum_{g_i \in \{f_i, \overline{f_i}\}} M_{g_{i_1}} \cdots M_{g_{i_k}} \\ &= \sum_{g_i \in \{f_i, \overline{f_i}\}} M_{g_{i_1} \circ g_{i_2} \circ \cdots \circ g_{i_k}}, \end{aligned}$$

where $g_{i_1} \circ g_{i_2} \circ \cdots g_{i_k}$ is the composed map which is monotone (possibly with an empty domain). This means that we can have at most 2^k monotone maps (and at least one since $P_{i_1} \cdots P_{i_k} \neq 0$). Therefore, the total number of maps is at most $d \cdot 2^k = d^{O(1)}$ as $k = O_{\lambda_0}(\log(1/\lambda))$. □

5.2.6 *Amplifying the Average Kazhdan Constant*

The *Kazhdan constant* is a notion of "spectral gap" (and so it is related to bias) for discrete groups, which predates (and was central to) the study of expansion in finite groups and graphs. In particular, we can work with finitely generated groups that can have infinitely many irreducible representations on more general Hilbert spaces, possibly of infinite dimension. Nonetheless, we can still apply our operator version of Ta-Shma's amplification procedure as it is independent of dimension and works for any unitary representation ρ. Therefore, we amplify the average Kazhdan constant, which also amplifies the Kazhdan constant. We now define these two parameters formally.

Let G be a group generated by a finite set S of generators. For a representation ρ, define

$$\begin{aligned}
\mathcal{K}(G,S,\rho) &:= \inf_{v\in\mathcal{H}\colon \|v\|_2=1} \max_{g\in S} \|\rho(g)\,v - v\|_2^2 \\
\overline{\mathcal{K}}(G,S,\rho) &:= \inf_{v\in\mathcal{H}\colon \|v\|_2=1} \frac{1}{|S|}\sum_{g\in S} \|\rho(g)\,v - v\|_2^2 \\
&= 2\left(1 - \left\|\mathop{\mathbb{E}}_{g\sim S}[\rho(g)]\right\|_{\mathrm{op}}\right).
\end{aligned}$$

Using these, we can define the Kazhdan constants by taking the minimum over representations,

$$\begin{aligned}
\mathcal{K}(G,S) &:= \inf\left\{\mathcal{K}(G,S,\rho) \mid (\rho,\mathcal{H}) \text{ irreducible and non-trivial unitary representation}\right\}. \\
\overline{\mathcal{K}}(G,S) &:= \inf\left\{\overline{\mathcal{K}}(G,S,\rho) \mid (\rho,\mathcal{H}) \text{ irreducible and non-trivial unitary representation}\right\}.
\end{aligned}$$

Here, the first definition is the usual definition of the Kazhdan constant of G with respect to generators S, whereas the second definition is an average version of the Kazhdan constant as in the work of Pak and Zuk [PZ02].

Corollary 5.2.19 (Amplifying Average Kazhdan Constant)**.** *Let* G *be a finitely generated group and* S *a finite set of generators such that the average Kazhdan constant* $\overline{\mathcal{K}}(G,S)$ *is equal to* $2\cdot(1-\lambda_0)$

for some constant $\lambda_0 \in (0,1)$*. For every* $\lambda \in (0,1)$*, there is a set* $S' \subseteq G$ *such that*

1. $\overline{\mathcal{K}}(G, S') \geq 2 \cdot (1-\lambda)$*, and thus,* $\mathcal{K}(G, S') \geq 2 \cdot (1-\lambda)$*.*
2. $|S'| = O_{\lambda_0}(|S|/\lambda^{2+o(1)})$*, and*
3. S' *can be found in time* $\mathrm{poly}(|S|/\lambda)$ *assuming an oracle for group operations on* G*.*

Remark 5.2.20. Note that the above amplification for $\overline{\mathcal{K}}$ immediately implies the same amplification for $\mathcal{K}$ (since the maximum is at least the average, $\overline{\mathcal{K}}(G,S) \leq \mathcal{K}(G,S)$) . Moreover, we remark that the above amplification can also similarly improve the constant of Lubotzky's property (τ) (the latter being a weaker version of property (T)), so it is more general and applies to expansion in many more discrete groups [RL10].

We will now apply this corollary to a specific family of representations, which will give a simple improvement to the bounds on the dimension expander constructed in [LZ08].

5.2.7 Explicit Almost Ramanujan Dimension Expanders

Dimension expanders were defined in [BISW01] motivated by applications in theoretical computer science. A conjectured construction based on irreducible representations was suggested by Wigderson to hold over every field. The conjecture was subsequently established by Lubotzky and Zelmanov [LZ08] for fields of characteristic zero. We now define dimension expanders, explain the [LZ08] proof, and our amplification in this setting.

Definition 5.2.21 ((ε,γ) Dimension Expander)**.** Let $\mathbb{F}$ be a field, $d \in \mathbb{N}$, $\varepsilon > 0$, V be a vector space of dimension n and $T_1, \ldots, T_d \colon V \to V$ be linear transformations. We say that $(V, \{T_i\}_{i\in[d]})$ is an (ε,γ)-dimension expander if for every subspace $W \subseteq V$ of dimension at most γn, we have $\dim(W + \sum_{i=1}^{d} T_i(W)) \geq (1+\varepsilon)\cdot \dim(W)$.

Remark 5.2.22. Observe that if the maps T_i are restricted to being permutation matrices, and the expansion condition is restricted only to subspaces W generated by elementary

basis vectors, one obtains the usual definition of vertex expansion of graphs. Thus, dimension expanders may be viewed as a linear-algebraic extension of expander graphs.

For an irreducible unitary representation ρ, there exists an associated representation[6] adj_ρ. The construction in [LZ08] relates dimension expansion with the Kazhdan constant. Their result gives a dimension expander, which gives expansion for all subspaces W, such that $\dim(W) \leq n/2$, but their expansion guarantee is significantly stronger when $\dim(W)$ is smaller. To obtain this, we first state a simple improvement to a computation in [LZ08].

Claim 5.2.23. *Let $W, W' \subseteq \mathbb{C}^d$ be two vector spaces. Let P, P' be orthogonal projectors onto W, W', respectively. Then,*

$$\mathrm{Re}\left(\mathrm{Tr}(PP')\right) = \mathrm{Tr}(PP') \geq \dim(W \cap W').$$

Proof. Let $\mathcal{U}_0 = \mathcal{W} \cap \mathcal{W}'$, $\mathcal{U}_1 = \mathcal{W} \cap \mathcal{U}_0^\perp$ and $\mathcal{U}_2 = \mathcal{W}' \cap \mathcal{U}_0^\perp$, with orthonormal bases $\{u_1, \ldots, u_k\}$, $\{a_1, \ldots, a_\ell\}$ and $\{b_1, \ldots, b_m\}$, respectively. We can write P and P' as

$$P = \sum_{i=1}^{k} u_i u_i^\dagger + \sum_{i=1}^{\ell} a_i a_i^\dagger \text{ and } P' = \sum_{i=1}^{k} u_i u_i^\dagger + \sum_{i=1}^{m} b_i b_i^\dagger.$$

Using linearity and orthogonality, we obtain

$$\begin{aligned}\mathrm{Tr}(PP') &= \sum_{i=1}^{k} \|u_i\|^2 \,\mathrm{Tr}(u_i u_i^\dagger) + \sum_{i=1}^{\ell}\sum_{j=1}^{m} \langle a_i, b_j\rangle \,\mathrm{Tr}(a_i b_j^\dagger) \\ &= k + \sum_{i=1}^{\ell}\sum_{j=1}^{m} \langle a_i, b_j\rangle^2 \geq \dim(\mathcal{U}_0),\end{aligned}$$

here in the last step we used that $\|u_i\|^2 = \mathrm{Tr}(u_i u_i^\dagger) = 1$. □

6. Let $\mathfrak{sl}_n(\mathbb{C}) = \{\mathrm{tr}(A) = 0 \mid A \in M_n(\mathbb{C})\}$. Equip the space with the Frobenius inner product defined as $\langle A, B\rangle = \mathrm{tr}(A^\dagger B)$ where $A^\dagger$ is the conjugate transpose. For any finite-dimensional unitary representation $\rho : G \to \mathbb{U}_n$, we have an adjoint representation $(\mathrm{adj}_\rho, \mathfrak{sl}_n)$ where the action is by conjugation $\mathrm{adj}_\rho(g) \cdot A = \rho(g) \cdot A \cdot \rho(g)^{-1}$. Since conjugation by unitary matrices preserves the trace, $\mathfrak{sl}_n$ is closed under the representation. Moreover, it is unitary as,

$$\left\langle \mathrm{adj}_\rho(g)A, \mathrm{adj}_\rho(g)B\right\rangle = \mathrm{tr}\left(\rho(g)A^\dagger\rho(g)^\dagger\rho(g)B\rho(g)^{-1}\right) = \langle A, B\rangle.$$

The above claim is a variant of the one used in [LZ08] to prove their main result. By plugging in Claim 5.2.23 in their proof we obtain,

Proposition 5.2.24 (Adapted from [LZ08] using Claim 5.2.23)**.** *Let* $\rho\colon G \to \mathbb{U}_{\mathbb{C}^n}$ *be a unitary irreducible representation. Then* $(\mathbb{C}^n, \{\rho(g)\}_{g\in S})$ *is* $(1-\lambda-o_n(1), 1/2 - O(\lambda))$*-dimension expander, where* $2(1-\lambda) := \mathcal{K}(G, S, adj_\rho)$*.*

Corollary 5.2.25. *Let* $\lambda > 0$ *be any fixed constant. Then, there exists an explicit infinite family of* $\left(1-\lambda-o_n(1), \frac{1}{2} - O(\lambda)\right)$*-dimension expanders.*

Proof. Pick a family of groups $\{G_n\}_n$ such that each G_n satisfies the condition of Corollary 5.2.19; for example, one can take any non-abelian finite simple group. By definition, for any such G, we have $\mathcal{K}(G, S, \mathrm{adj}_\rho) \geq \mathcal{K}(G, S)$ and therefore we obtain a set S' such that $\mathcal{K}(G, S', \mathrm{adj}_\rho) \geq 2(1-\lambda)$ for the given λ. We can now apply Proposition 5.2.24. □

Remark 5.2.26. Forbes and Guruswami [FG15] point out that the quantum expander construction of Harrow [Har07] also yields a dimension expander (with a similar construction of the dimension expanders from [LZ08]). As mentioned earlier, monotone expanders are dimension expanders over any field [DS09, DW10]. Moreover, the Bourgain and Yehudayoff [BY13] construction of monotone expanders with constant generating set yields such dimension expanders with a constant generating set.

5.2.8 Diameter of Finite Groups

The study of the diameter of Cayley graphs can take many forms, e.g., it can be with respect to every generating set (as in the celebrated Babai–Seress conjecture [BS88]) or with respect to some constant size generating set as in [BKL89]. Here, we explore the latter case.

First, recall that any n-vertex degree-d graph has diameter at least $\log_{d-1}(n)$. On the other hand, it is well-known that expansion directly implies diameter at most $C \cdot \log_{d-1}(n)$

for some constant $C \geq 1$ (depending on the expansion). The best upper bound a spectral proof can provide is 2 due to the Alon–Bopanna bound for spectral expansion. Using our amplification to almost-optimal spectral expansion, deduce that any expanding group G has a constant degree-d Cayley expander of diameter $2 + o_d(1)$.

Lemma 5.2.27. *Suppose $\{\mathrm{Cay}(G_i, S_i)\}_{i\in\mathbb{N}}$ is a family of bounded degree Cayley expanders. Then, there exists a family $\{\mathrm{Cay}(G_i, S_i')\}_{i\in\mathbb{N}}$ of constant degree-d Cayley expanders with diameter at most $(2 + o_d(1)) \cdot \log_{d-1}(G_i)$.*

Proof. We apply Theorem 1.2.2 to the family $\{\mathrm{Cay}(G_i, S_i)\}_{i\in\mathbb{N}}$ obtaining a new family of $\{\mathrm{Cay}(G_i, S_i')\}_{i\in\mathbb{N}}$ of (d, λ)-expanders with $d = 1/\lambda^{2+\beta}$ for some sufficiently small constants $\lambda, \beta > 0$. Let A_i be the normalized adjacency matrix of $\mathrm{Cay}(G_i, S_i')$ and $n_i = |G_i|$. Let e_g be the indicator vector of some fixed $g \in G_i$. Note that

$$\|(A_i - J/n_i)^t e_g\|_2 \;\leq\; \lambda^t = d^{-t/(2+\beta)} < 1/|G_i|,$$
$$\text{for } t = (2 + 2\beta) \cdot \log_d(|G_i|) \;=\; (2 + o_{d,\beta}(1)) \cdot \log_{d-1}(|G_i|).$$

This implies that $A_i^t e_g$ is supported on all elements of G_i, and thus the diameter of G_i is at most t. □

CHAPTER 6

DERANDOMIZED HOMOMORPHISM TESTING

Laqa, Piro, and Janki are out at the Promontory on a bright sunny day.

Piro: Pick one, a clever proof, or the "right definition".

Janki: I love an ingenious proof; I mean, there is no "Definitions from the book".

Laqa: Maybe we should write one, then. I would include Eymard's definition of the *algebra norm* from the 60s. I learned it while working on a problem, and it captured the underlying computation so beautifully that the proof wrote itself!

Piro: 'Algebra norm' rings a bell. Is it the trace norm of the convolution operator? Although I do not recall where it stems from.

Laqa: It is, but that definition I would not put in the book! It is a generalization of the spectral norm to functions on non-abelian groups.

Janki: Hmm, I do agree with you; one cannot overestimate the importance of Bassalygo and Pinsker defining expanders...

Laqa: ...or of Cayley defining groups abstractly and their "graphical representations".

In this chapter, we address Question 1.0.2 and show that small-bias sets fool functions with a small algebra norm. Using this, we show that small-bias sets can efficiently approximate Gowers' U^2-norm. This gives a randomness-efficient homomorphism test in the low-soundness regime. More generally, we prove that expanding Cayley graphs satisfy a "degree-2" *expander mixing lemma*. This structural result implies a derandomized version of the Babai–Nikolov–Pyber (BNP) lemma and the fooling of U^2-norm. Since the theorem statement is technical and the setup is different, we give a lengthier introduction.

6.1 Introduction

An important problem in theoretical computer science is efficiently testing if a function $f : G \to H$ correlates with some homomorphism between groups G and H. Such tests are widely used in constructions of *probabilistically checkable proofs* (PCPs), hardness of approximation, locally testable codes, and many other areas of computer science. Recently, there has been an interest in studying such tests for non-abelian groups. For example, in quantum complexity, *entanglement testing* [NV17] involves homomorphism testing over the (non-Abelian) Pauli group, which played an essential role in the proof of MIP*=RE [JNV$^+$21]. Additionally, such non-Abelian tests have been used to construct better PCPs [BK21] and for hardness of approximation results [BKM22].

The famous three-query randomized Blum–Luby–Rubinfield [BLR90] (BLR) test is as follows: pick two uniformly random group elements $x, y \in G$ and check if the homomorphism property holds for this pair, namely, if $f(x)f(y) = f(xy)$. This simple local test surprisingly sheds light on a global property of the function: if a function $f : \mathbb{Z}_2^n \to \mathbb{Z}_2$ passes the test with non-trivial probability, then the function must have a non-trivial correlation with some homomorphism.

This test can be used for any pair of groups G, H (assuming one can sample from G) and requires $2\log|G|$ random bits. A randomness-efficient version of the test that has been studied is the *derandomized BLR test* wherein x is uniformly sampled from G as before, but y is chosen from a sparse pseudorandom set $S \subseteq G$. If S is constant-sized, then the randomness is reduced to $\log|G| + O(1)$, which is almost optimal.

The study of such derandomized linearity (and low-degree) tests has found significant applications, particularly in the development of Probabilistically Checkable Proofs (PCPs). For example, the derandomization results in [BSVW03] have enabled the construction of PCPs and Locally Testable Codes of nearly linear size. Additionally, derandomized parallel execution [ST00, HW03] of the BLR test has facilitated the creation of PCPs with

low amortized costs. We extend this study of derandomized tests and investigate the question,

Given a function $f : G \to H$ *that passes the derandomized BLR test with probability* δ*, what can one conclude about the function* f*?*

An ideal conclusion would be that the function f is close to a true homomorphism φ in some metric, i.e., $\|f - \varphi\| \leq \theta(\delta)$. This is achievable in the "99%-regime", when the test passing probability, δ, is close to 1. This is also called the *unique-decoding regime*, as there is a unique homomorphism, φ, near the given function, f. The unique homomorphism can often be constructed via a *majority decoding procedure*. There are many results in this setting [BCH$^+$95, Far00, BP18] including derandomized ones [SW04]. In particular, for any finite group G and an arbitrary (not necessarily finite) group H, [Far00] constructs a homomorphism close in Hamming metric to the given function f if the test passing probability is $\geq 10/11$.

However, the situation is significantly more complex in the low-soundness or "1%-regime", wherein the function performs barely better than a random function, i.e., the test passing probability is $\frac{1}{|H|} + \delta$. Firstly, one cannot always hope to find a homomorphism close in the Hamming metric. A folklore counterexample due to Coppersmith (also in [BOCLR07]) gives a function $f : \mathbb{Z}_{3^k} \to \mathbb{Z}_{3^{k-1}}$ that passes the test with probability $2/9$ but it is far away from every homomorphism in the Hamming metric. More interestingly, for some pair of groups G, H, the only homomorphism from $G \to H$ might be the trivial one. In this case, we might not be able to conclude that f is close to the trivial homomorphism, but we can deduce something about the global structure of f. To do so, however, we need to better understand how the set of functions, from G to H, relates to the set of homomorphisms. For instance, Fourier analysis yields that any function $f : \mathbb{Z}_2^n \to \mathbb{C}$ can be expressed as a linear combination of homomorphisms. In general, *representation theory* gives a similar relation for the more general setting of functions $f : G \to \mathbb{U}_t$, where G is

any finite group, and $\mathbb{U}_t$ is the group of $t \times t$ unitary matrices. Therefore, this setting is a natural starting point for investigating the general question of derandomized testing for non-abelian groups. Moreover, this setting (which we work with throughout our paper) has exciting connections to quantum linearity testing!

6.1.1 Our Setup: Matrix-valued functions

We will work with functions from an arbitrary finite group G to the group of $t \times t$ unitary matrices, $f : G \to \mathbb{U}_t$, and will use the following inner product to measure correlation, i.e., $\langle f, g \rangle_{tr} = \mathbb{E}_{x \sim G}[\mathrm{tr}(g^*(x)f(x))]$. We wish to design a randomness-efficient variant of BLR such that if a function $f : G \to \mathbb{U}_t$ passes such a test, then the function f correlates with a homomorphism or a function arising from a homomorphism.

This setting has been studied in prior works [MR15, NV17, GH17], most notably in the context of quantum low-degree tests. The results of [MR15] and [GH17] are particularly relevant to our result, and we will discuss them in detail shortly. The other result by Natarajan and Vidick [NV17] gives a BLR-like test for homomorphism testing of functions, $f : \mathcal{P}^n \to \mathbb{U}_t$ where $\mathcal{P}^n$ is the n-fold tensor product of the *Pauli group* (also known as the *Weyl-Heisenberg* group). This was initially developed for entanglement testing [NV17] and later became a crucial component in the MIP* = RE proof [JNV+21].

While our setting encompasses their setup and has identical notions of correlation, our results do not directly apply to the quantum linearity test. This is because their test works with a specific presentation of the Pauli group due to additional constraints related to quantum measurements. Nevertheless, there might be exciting connections between our results and those in quantum homomorphism testing.

Before we state our results, we briefly define some relevant concepts and discuss the challenges associated with this setting. See Section 1.3 to recall the definitions.

Every finite group G has a finite set of *irreducible representations* (irreps) $\widehat{G}$, which are

the building blocks of complex-valued functions on G. In the case of Abelian groups, all the irreducible representations are one-dimensional and given by characters, also called Fourier characters. These characters also form an orthogonal basis for the space of complex-valued functions. For general finite groups, the orthogonal basis is given by the set of matrix coefficients of irreps, i.e., $\left\{\rho(x)_{i,j} \mid \rho \in \widehat{G}\ i,j \leq \dim(\rho)\right\}$. Just as in the Abelian case, we use $\hat{f}(\rho) = \mathbb{E}_x[f(x)\rho(x)]$ to represent the coefficient corresponding to irrep ρ, which is now a matrix. Such a basis also exists for matrix-valued functions, $f : G \to \mathbb{C}^{t\times t}$.

Challenges with this general setting

This general setting has three key differences from the original setting of BLR $f : \mathbb{Z}_2^n \to \mathbb{Z}_2$: (i) G is an arbitrary (not necessarily Abelian) finite group, (ii) H is continuous and not discrete, and (iii) the test passing probability is low (low-soundness regime). While each of these generalizations presents its own challenges, these are compounded when they are all together. In order to clearly illustrate the issues, let us focus only on the case of complex-valued functions, $f : G \to \mathbb{U}_1 \subseteq \mathbb{C}$. The entire discussion is relevant when the functions are matrix-valued, but this special case captures all the difficulties.

Hamming norm is unsuitable Since the codomain is continuous, the Hamming norm is inappropriate as it is sensitive to small perturbations. For example, let G be a finite group such that the only homomorphism to $\mathbb{U}_1$ is the trivial one. These groups exist and are known as *quasirandom groups*. If $f(x)$ is set to $e^{-i\varepsilon}$ for half of the inputs from G and $e^{-2i\varepsilon}$ for the rest, the BLR test passes with probability roughly $\frac{1}{8}$. But f has a normalized Hamming distance of 1 from the closest homomorphism, i.e., the trivial homomorphism. However, f is actually close to the trivial homomorphism in $\mathcal{L}^2$-distance, $\|f-g\|^2 = \mathbb{E}_x[|f(x)-g(x)|^2]$. This suggests why previous works [GH17, BFL03, MR15] in this setting have used the $\mathcal{L}^2$-norm.

Need to look at larger representations For abelian groups G, every function $f : G \to \mathbb{C}$

can be expressed as a linear combination of homomorphisms from $G \to \mathbb{C}$. This is no longer true when G is non-abelian. As we have seen in the preliminaries, we need to rely on homomorphisms $\rho : G \to \mathbb{U}_t$ for t potentially as large as $\sqrt{|G|}$, even though the original function maps to scalars. Thus, it is not immediate how to formalize the statement "f correlates with the homomorphism ρ," as $f : G \to \mathbb{C}$ whereas, $\rho : G \to \mathbb{U}_t$ for $t > 1$. There have been two non-equivalent solutions to this in prior work,

1. Clip[1] a representation – Let $g_\rho : G \to \mathbb{C}$ be defined as $g_\rho(x) = V\rho(x)U^*$ for a homomorphism ρ, and $V, U \in \mathbb{C}^{1\times t}$. One can now check if f correlates with g_ρ. This is the route taken by Gowers and Hatami [GH17].
2. Large Fourier Mass – If the representation ρ is irreducible, one can look at how much of the function can be explained by the Fourier basis elements corresponding to ρ, i.e., $\{\rho_{i,j}\}_{i,j}$. Note that these elements are no longer homomorphisms (unlike for Abelian groups), but $\|\hat{f}(\rho)\|_{\mathrm{HS}}$ being large is another way to formalize f being correlated with ρ, as $\|f\|^2 = \sum_\rho \dim(\rho)\|\hat{f}(\rho)\|^2_{\mathrm{HS}}$ by Parseval's. This approach is followed by [MR15].

Representation sizes depend on test-passing probability Ideally, we would want to show the correlation of the function with a representation ρ of dimension 1. But the above discussion shows why that is too much to ask for. Nevertheless, one might still want to bound the dimension of these representations, the intuition being that the smaller the dimension, the closer f is to being a true homomorphism. We will see that this can be done, but this dimension must depend on the test passing probability δ. Such a dependence is unavoidable, as the following counterexample illustrates. Let Γ be a non-abelian group containing an irrep of dimension $d \simeq \mathrm{poly}(|\Gamma|)$, say ρ. Consider

1. For functions $f : G \to \mathbb{U}_t$ for $t > 1$, the representations that need to be considered could also be of dimension $t' < t$. We still refer to it as "clipping".

the group $G = \mathbb{Z}_2^n \times \Gamma$ and its irrep, $\psi = \mathrm{triv} \otimes \rho$. The function $f(x) = \psi(x)_{1,1}$ passes the randomized BLR test if either of the query points is in the kernel of this irrep. Since $\mathbb{Z}_2^n$ lies in the kernel, the test passes with probability at least $\frac{1}{|\Gamma|}$. However, by construction, f is entirely supported on a d-dimensional irrep, ψ.

Known Results

As previously mentioned, despite the outlined challenges, researchers have achieved intriguing results in this matrix-valued function setting that we are interested in. Specifically, Moore and Russell [MR15] as well as Gowers and Hatmai [GH17] explored the homomorphism testing problem for functions from any finite group G to $\mathbb{U}_t$. It is important to note that their studies did not explicitly focus on this question from a testing perspective. Nevertheless, their findings can be easily adapted into a BLR-type linearity testing framework. Additionally, when translated into homomorphism testing terminology, their approaches correspond to the same Hilbert-Schmidt version of the BLR test that we employ. Below, we summarize the homomorphism testing results derived from these two studies.

Theorem 6.1.1 (Tests from [GH17, MR15])**.** *Let G be any finite group and $f : G \to \mathbb{U}_t$ be a unitary matrix-valued function. Assume that the function f passes the BLR test with probability δ. Then,*

1. *(Correlation with clipped representation* [GH17]*) : There is a representation, $\pi : G \to \mathbb{U}_{t'}$ and two matrices, $V, U \in \mathbb{C}^{t \times t'}$, such that f correlates with the function, $g_\pi = V\pi(x)U^* : G \to \mathbb{U}_t$:*

$$\left\langle f, g_\pi \right\rangle_{\mathrm{tr}} \geq \frac{\delta^2}{4} \cdot t$$

   Moreover, the representation has bounded dimension, $\delta^2 t \leq t' = \dim(\pi) \leq \frac{2t}{\delta^2}$.

2. *(Fourier Mass on a low-dimensional irreducible representation* [MR15]*) : There exists an irrep $\rho \in \widehat{G}$ such that $\dim(\rho) < \frac{2t}{\delta^2}$ and $\|\widehat{f}(\rho)\|_{\mathrm{HS}}^2 \geq \frac{\delta^2}{2}$.*

Remark 6.1.2. Note that the representation, π, in part 1 of the theorem is not guaranteed to be irreducible, but the second one is.

6.1.2 Our Results

The main contribution of this work is to give a derandomized BLR-like homomorphism test in the low soundness regime for the general setup of functions from an arbitrary finite group $G \to \mathbb{U}_t$. Prior to this work, the only known derandomized test in the 1%-regime is that of [BSVW03] for the case when $G = \mathbb{Z}_p^n$ and $H = \mathbb{Z}_p$. Our key derandomization tool is *small–bias sets* (Definition 1.3.10). We first review a few constructions.

Theorem 6.1.3 ([WX08, Thm 5.1])**.** *For every finite group* G *and any constant* $\varepsilon > 0$*, there exists a deterministic* $\mathrm{poly}(|G|)$*-time algorithm that outputs an* ε*-biased set* $S \subseteq G$ *of size* $|S| \leq O\left(\frac{\log |G|}{\varepsilon^2}\right)$.

While this bound is tight over Abelian groups, we can do much better for other groups. Particularly for all *finite simple groups*, we now have explicit constant-sized small-biased sets due to a long line of work [KN06, Kas07, Lub11]. These can also be made near-optimal using the amplification machinery from Chapter 4, Theorem 5.2.1.

Theorem 6.1.4 ([KN06, Kas07, Lub11], Theorem 5.2.1)**.** *For every non-abelian finite simple group* G *and any* $\varepsilon > 0$*, there exists a deterministic* $\mathrm{poly}(1/\varepsilon)$*-time algorithm that outputs an* ε*-biased set* $S \subseteq G$ *of size* $|S| \leq O\left(\varepsilon^{-(2+o(1))}\right)$.

Result 1: Derandomized Homomorphism Testing We analyze the following derandomized variant of BLR. Here, the parameter γ allows for a relaxed version of this test that makes the test robust to small noise, which can be useful as $\mathbb{U}_t$ is a continuous group.

Derandomized BLR$_\gamma$ (G, S, f):

1. Sample $x \sim G, y \sim S$.
2. If $\|f(xy) - f(x) \cdot f(y)\|_{HS}^2 \leq \gamma t$, output *Pass*. Else, output *Fail*.

Setting $\gamma = 0$ recovers the usual derandomized version of the BLR test used in previous derandomizations of homomorphism tests [BSVW03, SW04].

Theorem 6.1.5 (Informal version of Theorem 6.4.4)**.** *Let G be any finite group and $f : G \to \mathbb{U}_t$ be a unitary matrix-valued function. Let $S \subseteq G$ be an ε-biased set. Assume that the function f passes the derandomized BLR test with probability $\delta > \sqrt{\varepsilon}$. Then,*

1. *(Correlation with clipped representation): There is a representation, $\pi : G \to \mathbb{U}_{t'}$ and two matrices, $V, U \in \mathbb{C}^{t \times t'}$, such that for $g_\pi = V\pi(x)U^* : G \to \mathbb{U}_t$, f correlates with g_π,*

$$\left\langle f, g_\pi \right\rangle_{tr} \geq \frac{\delta^2 - \varepsilon}{4} \cdot t$$

   Moreover, the representation has bounded dimension, $(\delta^2 - \varepsilon)t \leq t' = \dim(\pi) \leq \frac{2t}{\delta^2 - \varepsilon}$.

2. *(Fourier Mass on a low-dimensional irreducible representation): There exists an irrep $\rho \in \widehat{G}$ such that $\dim(\rho) < \frac{2t}{\delta^2 - \varepsilon}$ and $\|\widehat{f}(\rho)\|_{HS}^2 \geq \frac{\delta^2 - \varepsilon}{2}$.*

Moreover, if one uses the γ-robust BLR test, the same conclusions hold with δ replaced by $\delta - (\gamma/2)$.

Using the small-bias set construction from Theorem 6.1.3, we get a test that uses $\log|G| + \log|S| = \log|G| + O(\log\log|G|) = (1 + o(1))\log|G|$-random bits. For special families of groups like the class of *finite simple groups*, we can use Theorem 6.1.4 to further reduce the randomness to $\log|G| + O(1)$, which is almost optimal.

Result 2: Derandomized BNP Lemma The "BNP lemma" is a very useful observation due to Babai, Nikolov, Pyber [BNP08], and Gowers [Gow08]. This lemma gives

an improvement over Cauchy–Schwarz for *quasirandom groups*, i.e., groups with no small non-trivial irreps, and has been used to analyze mixing in progressions [BHR22], product-free sets [Gow08], and hardness of approximation [BKM22], to name a few. In its most general form, it says that for functions to $t \times t$-matrices, $f, g : G \to M_t(\mathbb{C})$, we have,

$$\|f * g\|^2 = \mathop{\mathbb{E}}_{s\sim G}\left[\|(f*g)(s)\|_{HS}^2\right] \;\le\; \frac{1}{D}\,\|f\|_2^2\|g\|_2^2.$$

We show that such a bound holds even when the average is over a small-bias set $S \subseteq G$, which could be of constant size for some groups!

Lemma 6.1.6 (Derandomized Matrix BNP)**.** *Let G be a group such that the dimension of the smallest non-trivial irrep is D, and let $S \subseteq G$ be an ε-biased set. Let $f, g : G \to M_t(\mathbb{C})$ be mean-zero functions. Then,*

$$\mathop{\mathbb{E}}_{s\sim S}\left[\,\|(f*g)(s)\|_{HS}^2\right] \;\le\; \left(\frac{1}{D}+\varepsilon\right)\|f\|_2^2\|g\|_2^2$$

The usual BNP lemma can be recovered by setting $S = G$, and thus, $\varepsilon = 0$.

6.1.3 Technical Overview

Our main conceptual contribution is to initiate the study of the non-abelian generalization of two useful notions in the analysis of Boolean functions: (i) *spectral norm* of a function and (ii) *spectral positivity*.

Spectral norm and its non-abelian analog The ℓ_1-norm of the Fourier transform of a function is known as its *spectral norm*. Spectral norm has emerged as an important quantity for the analysis of Boolean functions, i.e., functions over $\mathbb{Z}_2^n$. In particular, functions with low spectral norm have a lot of structure [STV17]: they admit small decision trees, parity decision trees, they are easily learnable, etc.

One of the conceptual contributions of this paper is studying the non-abelian analog

of this norm from the perspective of pseudorandomness. A first generalization one can think of would be a similar ℓ_1 norm of the Fourier coefficients. However, it turns out that the appropriate generalization of the spectral norm is the *Fourier algebra norm*. This was suggested earlier by Sanders [San21], who used it to generalize the quantitative idempotent theorem. This norm has multiple equivalent definitions, but our key idea is to use the following harmonic analytic reformulation due to Sanders [San21] (attributed to [Eym64]),

$$\|f\|_A = \min_{(\pi,V)} \left\{ \|u\| \cdot \|v\| \,\middle|\, f(x) = \langle u, \pi(x)\, v\rangle \right\}$$

where (π, V) is a representation of G and $u, v \in V$. [2]

It is well-known that any function, f, on an Abelian group is $\varepsilon\|\hat{f}\|_1$-fooled by any ε-biased set. We show that this neatly generalizes to any finite group by replacing the spectral norm with Fourier algebra norm, any function, f, on a finite group is $\varepsilon\|f\|_A$-fooled by any ε-biased set.

Spectral Positivity and its non-abelian analog A function over an Abelian group G, $f : G \to \mathbb{C}$, is spectrally non-negative of $\hat{f}(\chi) \geq 0$ for every character χ. This notion played a key role in the recent breakthrough by Kelley and Meka [KM23] on 3-AP free sets.

This naturally generalizes to the finite group setting wherein a *positive-definite functions* is a function f such that $\hat{f}(\rho)$ is positive semi-definite for every irreducible representation ρ. The critical observation is that such functions have a small algebra norm, $\|f\|_A = f(1)$. We use this to prove that small–bias sets can be used to approximate the U^2-norm.

2. The Fourier inversion theorem gives one such an expression for f by using the *regular representation;*. However, it might not be the one that minimizes the algebra norm; hence, one minimizes over such expressions.

Proof Overview

Denote $\tilde{f}(x) = f(x^{-1})^*$, and recall the following two norms,

$$\text{(U^2-norm)} \quad \|f\|_{U^2} = \|f * \tilde{f}\|^2 = \mathop{\mathbb{E}}_{x\sim G}\left[\|(f * \tilde{f})(x)\|_{HS}^2\right]$$

$$\text{(Algebra norm)} \quad \|f\|_A = \min_{(\pi,V)}\left\{ \|u\| \cdot \|v\| \;\middle|\; f(x) = \langle u, \pi(x)\, v\rangle \right\}$$

We now give a quick summary of the key steps involved in the proof:

1. (Lemma 6.3.1) Small–bias sets fool functions with a small algebra norm.

2. (Lemma 6.5.1) Let $f, g : G \to \mathbb{U}_t$ be any functions. Then, the function, $x \mapsto \|(f * g)(x)\|_{HS}^2$ has a small Fourier algebra norm.

3. The above two lemmas imply a degree-2 EML. This immediately yields our result on the derandomized BNP lemma (Lemma 6.1.6). We expect that this degree-2 EML will have uses beyond this work, and we explain this below.

4. A special case of the above EML implies that small bias sets approximate U^2-norm (Corollary 6.1.7). Thus, the derandomized test's passing probability implies a large U^2-norm of the function. Combining this with the inverse theorem of Gowers-Hatami [GH17] gives us the first part of Theorem 6.1.5.

5. The second part of Theorem 6.1.5 follows from the same large U^2-norm consequence implied by test passing. To achieve this, we adapt the proof strategy of the BNP lemma [BNP08] to our setup, which relies on basic non-abelian harmonic analysis.

Degree-2 EML Our key technical contribution is a degree-2 variant of the celebrated *expander mixing lemma* (EML). Recall that EML characterizes spectral expansion. When applied to the Cayley graph $\mathrm{Cay}(G, S)$, we get that S is an ε-biased set if and only if the

EML holds,

$$\left| \mathop{\mathbb{E}}_{s\sim S}[(f*g)(s)] - \mathop{\mathbb{E}}_{s\sim G}[(f*g)(s)] \right| \;\le\; \varepsilon \|f\|_2 \, \|g\|_2, \quad \text{(EML)}\,.$$

We prove that such sets also satisfy a degree-2 variant of the above inequality,

$$\left| \mathop{\mathbb{E}}_{s\sim S}\left[\|(f*g)(s)\|^2\right] - \mathop{\mathbb{E}}_{s\sim G}\left[\|(f*g)(s)\|^2\right] \right| \;\le\; \varepsilon \|f\|_2^2 \, \|g\|_2^2, \quad \text{(Our degree-2 EML)}\,.$$

Using it for the special case of where $g(x) = \tilde{f}(x) = f(x^{-1})^*$, we get that small bias sets approximate U^2-norm.

Corollary 6.1.7 (Small bias sets approximate U^2-norm)**.** *For $f : G \to \mathbb{U}_t$ such that $\|f\| = 1$,*

$$\left| \mathop{\mathbb{E}}_{s\sim S}\left[\|(f*\tilde{f})(s)\|^2\right] - \|f\|_{U^2} \right| \;\le\; \varepsilon \|f\|_2^4\,.$$

Thus, the U^2-norm of a function f can be ε-estimated by querying $f*\tilde{f}$ on an ε-biased set S.

6.1.4 Related Work

High soundness regime Blum–Luby–Rubinfield [BLR90] analyzed linearity tests for functions of the form $f : \mathbb{Z}_2^n \to \{\pm 1\}$. This was extended to the setting $f : G \to H$ where both are arbitrary finite groups, by Ben-Or, Coppersmith, Luby, and Rubinfeld [BOCLR07]. This result was derandomized by Shpilka and Wigderson [SW04]. Going beyond finite groups, Farah [Far00], and later, Badora and Przebieracz [BP18], give homomorphism tests for any *amenable group* G, and any group H, equipped with an invariant metric.

Low soundness regime Bellare, Coppersmith, Håstad, Kiwi, and Sudan [BCH+95] analyzed linearity tests for functions of the form $f : \mathbb{Z}_2^n \to \{\pm 1\}$ in this low-soundness regime. This was extended to the setting, $f : \mathbb{Z}_p^n \to \mathbb{Z}_p$, by Håstad and Wigderson [HW03]. This result was derandomized using ε-biased sets by Ben-Sasson, Sudan, Vadhan, and Wigderson [BSVW03]. For the same setting, Kiwi [Kiw03] analyzed a variant of the BLR test that

uses a lot more randomness but gives an improved correlation. Samorodnitsky [Sam07] studied a completely different setup where H is large and not a subset of $\mathbb{C}$. He showed that if a function $f : \mathbb{Z}_2^n \to \mathbb{Z}_2^m$ passes the test with probability δ, then it has an exponentially small agreement with a homomorphism. Improving the agreement to polynomial in δ is equivalent to the polynomial Freiman–Rusza (PFR) conjecture which was finally settled recently [San12, GGMT23]. The following table summarizes previous work.

Work	Setting ($f : G \to H$)	Conclusion	Randomness
	High Soundness		
[BLR90]	$G = \mathbb{Z}_2^n$, $H = \mathbb{Z}_2$	Hamming	$2\log\lvert G\rvert$
[BOCLR07]	G, H any finite groups	Hamming	$2\log\lvert G\rvert$
[SW04]	G, H any finite groups	Hamming	$(1+o(1))\log\lvert G\rvert$
	Low Soundness		
[BCH+95]	$G = \mathbb{Z}_2^n$, $H = \mathbb{Z}_2$	Hamming	$2\log\lvert G\rvert$
[Kiw03]	$G = \mathbb{Z}_p^n$, $H = \mathbb{Z}_p$	Hamming	$(2+o(1))\log\lvert G\rvert$
[BSVW03]	$G = \mathbb{Z}_p^n$, $H = \mathbb{Z}_p$	Hamming	$(1+o(1))\log\lvert G\rvert$
[Sam07, San12, GGMT23]	$G = \mathbb{Z}_2^n$, $H = \mathbb{Z}_2^m$	Hamming	$2\log\lvert G\rvert$
[BFL03]	G finite abelian, $H = \mathbb{U}_1$	Correlation	$2\log\lvert G\rvert$
[MR15]	G any finite group, $H = \mathbb{U}_t$	Correlation	$2\log\lvert G\rvert$
[GH17]	G any finite group, $H = \mathbb{U}_t$	Hilbert-Schmidt	$2\log\lvert G\rvert$
Our Result	G any finite group, $H = \mathbb{U}_t$	Hilbert-Schmidt	$(1+o(1))\log\lvert G\rvert$
Our Result	G any finite group, $H = \mathbb{U}_t$	Correlation	$(1+o(1))\log\lvert G\rvert$

Table 6.1: Summary of prior works on homomorphism testing

6.2 Prelims: Matrix–valued functions and U^2-norm

Denote by $\mathcal{L}_t^2(G) = \{f : G \to M_t(\mathbb{C})\}$, the space of $t \times t$ matrix-valued functions equipped with the trace expectation inner product,

$$\langle f, g\rangle = \mathop{\mathbb{E}}_{x\sim G}[\langle f(x), g(x)\rangle_{tr}] = \mathop{\mathbb{E}}_{x\sim G}\left[\mathrm{Tr}\left(g(x)^* f(x)\right)\right] \tag{6.1}$$

The induced norm is $\|f\|^2 = \mathbb{E}_{x\sim G}[\|f(x)\|_{HS}^2]$. For a function f, we denote its *adjoint* by $\tilde{f}(x) := f(x^{-1})^*$. The operation of convolution generalizes as,

$$(f * g)(x) := \mathbb{E}_{y\sim G}\left[f(xy^{-1})g(y)\right] = \mathbb{E}_{y\sim G}\left[\tilde{f}(y)^* g(yx)\right].$$

Definition 6.2.1 (Matrix Fourier Coefficient)**.** For any irrep ρ, we have $\widehat{f}(\rho) := \mathbb{E}_x\left[f(x) \otimes \rho(x)\right]$. We denote the coefficient of the trivial irrep as $\mu(f) := \widehat{f}(\rho_{\text{triv}})$.

Fact 6.2.2. *The following identities hold for the matrix Fourier transform,*

1. **(Parseval's identity)** $\|f\|^2 = \mathbb{E}_x\left[\|f(x)\|_{\text{HS}}^2\right] = \sum_{\rho\in\widehat{G}} d_\rho \|\widehat{f}(\rho)\|_{\text{HS}}^2$
2. **(Convolution identity)** $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$
3. **(U^2 norm)** $\|f\|_{U^2} = \|\tilde{f} * f\|^2 = \sum_\rho d_\rho \left\|\widehat{f}(\rho)\widehat{f}(\rho)^*\right\|_{\text{HS}}^2 = \sum_\rho d_\rho \left\|\widehat{f}(\rho)^*\widehat{f}(\rho)\right\|_{\text{HS}}^2$

Proof. These facts are simple extensions of the scalar-valued functions and were considered in [BFL03, MR15, GH17]. Since the last one is perhaps atypical, we provide a quick proof.

$$\begin{aligned}
\|f\|_{U^2} &:= \mathop{\mathbb{E}}_{xy^{-1}=wz^{-1}}\left[\operatorname{Tr}\left(f(x)f(y)^* f(z)f(w)^*\right)\right] \\
&= \mathop{\mathbb{E}}_{t=xy^{-1}=wz^{-1}}\left[\langle f(x)f(y)^*, f(w)f(z)^*\rangle\right] \\
&= \mathop{\mathbb{E}}_{t=xy^{-1}=wz^{-1}}\left[\left\langle f(x)\tilde{f}(y^{-1}), f(w)\tilde{f}(z^{-1})\right\rangle\right] \\
&= \mathop{\mathbb{E}}_{t}\left[\left\langle (f*\tilde{f})(t), (f*\tilde{f})(t)\right\rangle\right] = \|f*\tilde{f}\|^2 \ .
\end{aligned}$$

The second equality follows from Parseval's identity and convolution identity once one observes that,

$$\widehat{\tilde{f}}(\rho) = \mathop{\mathbb{E}}_{x}\left[f(x^{-1})^* \otimes \rho(x)\right] = \mathop{\mathbb{E}}_{x}\left[f(x)^* \otimes \rho(x^{-1})\right] = \mathop{\mathbb{E}}_{x}[f(x)^* \otimes \rho(x)^*] = \hat{f}(\rho)^* \qquad \square$$

6.2.1 Fourier algebra norm and positive definite functions

Let $f : G \to \mathbb{C}$ be any function and let T_f be the convolution by f operator, $T_f(g) = f * g$. More explicitly, $T_f(x, y) = \frac{1}{|G|} f(x^{-1}y)$ is a $|G| \times |G|$ matrix.

Definition 6.2.3 (Fourier algebra norm)**.** The *algebra norm* of $f : G \to \mathbb{C}$ has the following equivalent definitions,

i. $\|f\|_A = \sup\{\langle f, g\rangle \mid \|T_g\|_{op} \leq 1\}$.

ii. $\|f\|_A = \|T_f\|_{tr} = \sum_i \sigma_i(T_f)$ where $\{\sigma_i\}$ are the singular values of T_f.

iii. $\|f\|_A = \min_{(\pi,V)}\{\|u\| \cdot \|v\| \mid f(x) = \langle u, \pi(x)\, v\rangle\}$ where (π, V) is a representation of G and $u, v \in V$.

The equivalence of definitions (i) and (ii) can be found in [San11, Lem. 5.2] and also in [HHH22, Prop 3.11]. The proof of equivalence between (i) and (iii) is present in [San21, Lem. 2.2], attributed to Eymard [Eym64, Thèorém, Pg 218]. In particular, Sanders shows that the minimum is indeed attained for some representation (π, V).

Abelian Case When the group is abelian, the algebra norm coincides with the *spectral norm*, i.e., $\|f\|_A = \|\hat{f}\|_1$. This can be seen by observing that T_f is a diagonal matrix (in the Fourier basis) with the Fourier coefficients on the diagonal. Therefore, the algebra norm generalizes the spectral norm to the non-Abelian setting.

Definition 6.2.4 (Positive definite functions)**.** Let G be a finite group. A function $f : G \to \mathbb{C}$ is said to be positive definite if the convolution operator, T_f, is positive semi-definite.

The following simple observation states that positive definite functions have a small algebra norm. We will crucially use this later to bound the algebra norm of a function.

Observation 6.2.5. If a function f is positive-definite, then $\|f\|_A = f(1)$.

Proof. Since T_f is positive semi-definite, $\|f\|_A = \|T_f\|_{tr} = \mathrm{tr}(T_f) = f(1)$ as $T_f(x, y) = \frac{f(x^{-1}y)}{|G|}$. □

6.3 Small-bias sets "fool" small norm functions

In the Boolean setting, properties of low spectral norm function have been well-studied [KM93, STV17]. Green and Sanders [GS08] showed that functions with low spectral norm can be expressed as a ± 1 combination of characteristic functions of cosets. Sanders [San11] generalized it to non-abelian groups with spectral norm replaced by the algebra norm.

This suggests that the algebra norm is indeed the suitable generalization of the spectral norm, and one might investigate other properties of functions with low algebra/spectral norm. This section makes another connection by showing that small-bias sets fool functions with small algebra norms, again generalizing the Abelian case.

Boolean Cube Let $f : \mathbb{Z}_2^n \to \mathbb{C}$, be any function and let S be an ε-biased set. Then,

$$\begin{aligned}\left|\mathop{\mathbb{E}}_{h\sim S}[f(h)] - \mathop{\mathbb{E}}_{h\sim G}[f(h)]\right| &= \left|\sum_{\chi}\hat{f}(\chi)\Big(\mathop{\mathbb{E}}_{h\sim S}[\chi(h)] - \mathop{\mathbb{E}}_{h\sim G}[\chi(h)]\Big)\right| \\ &\leq \max_{\chi}\left|\mathop{\mathbb{E}}_{h\sim S}[\chi(h)] - \mathop{\mathbb{E}}_{h\sim G}[\chi(h)]\right| \cdot \sum_{\chi\neq 0}\left|\hat{f}(\chi)\right| \\ &\leq \varepsilon \cdot \|f - \mu(f)\|_A \ .\end{aligned}$$

Moreover, this is tight due to a result of [DETT10, Prop. 2.7], which says that any function that is fooled by all ε-biased sets must be sandwiched by low spectral norm functions.

General finite groups We now generalize the above result for finite groups. The key here is to use the harmonic analytic definition of the spectral norm, which makes the proof surprisingly simple.

Lemma 6.3.1. *Let $f : G \to \mathbb{C}$, be a function and $S \subseteq G$ be any ε-biased set . Then,*

$$\left|\mathop{\mathbb{E}}_{x\sim S}[f(x)] - \mathop{\mathbb{E}}_{x\sim G}[f(x)]\right| \leq \varepsilon \cdot \|f\|_A \ .$$

Proof. From Definition 6.2.3, we get that there is a representation, (π, V), representation of G such that $f(x) = \langle u, \pi(x)\, v\rangle$ for some $u, v \in V$. Moreover, $\|u\|\|v\| \leq \|f\|_A$. Using Maschke's theorem (Theorem 1.3.8), we have $\pi = \rho_{\text{triv}}^{\oplus c} \oplus_i \rho_i$ where c denotes the multiplicity of the trivial representation and ρ_i are all non-trivial irreducible representations (possibly with repetitions). Let U_π be the unitary transformation that block-diagonalizes π. Then, $U_\pi v = v_{\text{triv}} \oplus v_i$ and similarly $U_\pi u = u_{\text{triv}} \oplus u_i$. Thus, we have,

$$\begin{aligned}\langle u, \pi(x)\, v\rangle = \langle U_\pi u, U_\pi\, \pi(x)\, v\rangle &= \langle U_\pi u, (U_\pi \pi(x) U_\pi^*)\, U_\pi v\rangle \\ &= \langle u_{\text{triv}} \oplus_i u_i, v_{\text{triv}} \oplus_i (\rho_i v_i)\rangle \\ &= \langle u_{\text{triv}}, v_{\text{triv}}\rangle + \sum_i \langle u_i, \rho_i(x)\, v_i\rangle\end{aligned}$$

Now, $\langle u_{\text{triv}}, v_{\text{triv}}\rangle$ is a constant and is the same for both terms. Therefore,

$$\begin{aligned}\mathop{\mathbb{E}}_{x\sim S}[f(x)] - \mathop{\mathbb{E}}_{x\sim G}[f(x)] &= \sum_i \left\langle u_i, \left(\mathop{\mathbb{E}}_{x\sim S}[\rho_i(x)] - \mathop{\mathbb{E}}_{x\sim G}[\rho_i(x)]\right) v_i\right\rangle \\ \left|\mathop{\mathbb{E}}_{x\sim S}[f(x)] - \mathop{\mathbb{E}}_{x\sim G}[f(x)]\right| &\leq \sum_i \|u_i\|_2 \cdot \left\|\mathop{\mathbb{E}}_{x\sim S}[\rho_i(x)] - \mathop{\mathbb{E}}_{x\sim G}[\rho_i(x)]\right\|_{\text{op}} \|v_i\|_2 \quad \text{(Cauchy–Schwarz)} \\ &\leq \varepsilon \sum_i \|u_i\|_2 \|v_i\|_2 \quad \text{(S is an ε-bias set)} \\ &\leq \varepsilon \cdot \sqrt{\sum_i \|u_i\|_2^2}\sqrt{\sum_i \|v_i\|_2^2} \quad \text{(Cauchy–Schwarz)} \\ &= \varepsilon \cdot \|u - u_{\text{triv}}\|_2 \|v - v_{\text{triv}}\|_2 \quad \text{(U_π is unitary)}.\end{aligned}$$

□

Corollary 6.3.2 (PD functions are fooled)**.** *If $f : G \to \mathbb{C}$ is a positive-definite function, then $\|f\|_A = f(1) \leq \|f\|_\infty$. Therefore, if $\|f\|_\infty \leq 1$, then f is ε-fooled by every ε-biased set.*

Proof. Since, f is positive definite, T_f is PSD and thus, $\|f\|_A = \|T_f\|_{\text{tr}} = \text{Tr}(T_f) = f(1) \leq 1$. □

6.3.1 U^2 norm and algebra norm

Let $f : G \to M_t(\mathbb{C})$ be a function, then $\|f\|_{U^2} = \mathbb{E}_{y\sim G}[\psi(y)]$ where,

$$\psi(y) = \left\|(\tilde{f} * f)(y)\right\|_{HS}^2 = \left\| \mathop{\mathbb{E}}_{x\sim G}[f(x)^* f(xy)] \right\|_{HS}^2 .$$

Therefore, if ψ has small algebra norm then, $\|f\|_{U^2}$ can be approximated by averaging ψ over an ε-biased set. We prove the algebra norm bound by proving that the function, ψ, is a positive-definite function.

Lemma 6.3.3. *For $f : G \to M_t(\mathbb{C})$, the function $\psi(y) = \left\|(\tilde{f} * f)(y)\right\|_{HS}^2$ is positive-definite.*

Proof. To prove that T_ψ is a PSD matrix, we wish to show that for any $\{c_a\}_{a\in G} \in \mathbb{C}^G$,

$$\sum_{a,b\in G} c_a \overline{c_b}\, \psi(a^{-1}b) \;\geq\; 0$$

The key observation is that, if we can write $\psi(a^{-1}b) = \mathbb{E}_{x,y\sim G}[\langle N_{x,y}(a), N_{x,y}(b)\rangle_{tr}]$ for some $N_{x,y}$, then ψ is positive-definite. We first show such a factorization of $\psi(a^{-1}b)$.

$$\begin{aligned}
\psi(a^{-1}b) &= \left\| \mathop{\mathbb{E}}_{x\sim G}\left[f(x)^* f(xa^{-1}b)\right] \right\|_{HS}^2 \\
&= \left\| \mathop{\mathbb{E}}_{x\sim G}[f(xa)^* f(xb)] \right\|_{HS}^2 \\
&= \left\langle \mathop{\mathbb{E}}_{x\sim G}[f(xa)^* f(xb)], \mathop{\mathbb{E}}_{y\sim G}[f(ya)^* f(yb)] \right\rangle_{tr} \\
&= \mathop{\mathbb{E}}_{x,y\sim G}[\langle f(xa)^* f(xb), f(ya)^* f(yb)\rangle_{tr}] \\
&= \mathop{\mathbb{E}}_{x,y\sim G}[\langle f(ya) f(xa)^*, f(yb) f(xb)^*\rangle_{tr}] \\
&:= \mathop{\mathbb{E}}_{x,y\sim G}\left[\langle N_{x,y}(a), N_{x,y}(b)\rangle_{tr}\right]
\end{aligned}$$

The second last equality uses the cyclicity of trace. The result now follows as,

$$\begin{aligned}
\sum_{a,b\in G} c_a \overline{c_b}\, \psi(a^{-1}b) &= \sum_{a,b\in G} c_a \overline{c_b} \mathop{\mathbb{E}}_{x,y\sim G}\left[\langle N_{x,y}(a), N_{x,y}(b)\rangle_{tr}\right] \\
&= \mathop{\mathbb{E}}_{x,y\sim G}\left[\left\langle \sum_{a\sim G} c_a N_{x,y}(a), \sum_{b\sim G} c_b N_{x,y}(b)\right\rangle_{tr}\right] \\
&= \mathop{\mathbb{E}}_{x,y\sim G}\left[\left\| \sum_{a\sim G} c_a N_{x,y}(a)\right\|_{HS}^2\right] \geq 0.
\end{aligned}$$

□

Remark 6.3.4. One can also deduce this by coupling *Stinespring's dilation theorem* with the observation in [DCOT18] that $\tilde{f} * f$ is *completely positive*. This immediately yields that $\psi(y) = \langle VV^*, \pi(y)VV^*\pi(y)^* \rangle = \langle VV^*, \rho(y)VV^* \rangle$ for some representation ρ. We will use this idea to prove Lemma 6.5.1, a general version of the above lemma.

6.4 Derandomized Matrix Correlation Testing

In this section, we will focus on functions of the form $f : G \to \mathbb{U}_t$. Let $S \subseteq G$ be an ε-biased set. We consider the following robust variant of the BLR test on group G:

$\mathbf{BLR}_\gamma(G, S, f)$:

1. Sample $x \sim G, y \sim S$.
2. If $\|f(xy) - f(x) \cdot f(y)\|_{HS}^2 \leq \gamma t$, output *Pass*. Else, output *Fail*.

If $S = G$, i.e., in the full randomness regime, it can be easily shown that if a function passes the test, it must have a large U^2-norm. Our key technical claim (Claim 6.4.2) is that if S is a small-biased set, then, essentially, the same conclusion can be drawn from derandomized BLR test passing.

This lower bound on the U^2-norm can then be plugged into the result of Gowers and Hatami [GH17], who showed that if a matrix-valued function on a finite group has non-trivial U^2-norm, then it must be close to some genuine representation. More specifically,

Theorem 6.4.1 (Gowers–Hatami [GH17])**.** *Let G be any finite group and let $f : G \to M_t(\mathbb{C})$ be a matrix-valued function such that $\|f(x)\|_{op} \leq 1$ and $\|f\|_{U^2} \geq ct$, for some $c > 0$. Then there are $t' \in [\frac{c}{2-c}t, \frac{2-c}{c}t]$ and a function $g(x) := V\pi(x)U^*$ where π is a t' dimensional unitary representation, U, V are $t \times t'$ dimensional partial unitary matrices, such that:*

$$\mathbb{E}_{x\sim G}\Big[\big\langle f(x), g(x)\big\rangle_{\mathrm{HS}}\Big] \;\geq\; c^2/4 \;.$$

We first prove the central derandomization claim, which lets us move from the test passing probability of the derandomized test, to a claim about the U^2 norm over the entire group.

Claim 6.4.2 (Derandomized Test also implies large U^2-norm)**.** *Let* $\gamma, \delta \geq 0$ *and* $f : G \to \mathbb{U}_t$*. If* f *passes the* $BLR_{2\gamma}(G, S, f)$ *test with probability* $\geq \delta$ *then,*

$$\|f\|_{\mathrm{U}^2} \;\geq\; \big((\delta-\gamma)^2 - \varepsilon\big)\cdot t.$$

Proof. Let, $\Delta(x,y) := \|f(x)f(y) - f(xy)\|_{\mathrm{HS}}^2$ and $\delta' = \delta - \gamma$. We have,

$$\mathbb{E}_{x\sim G, y\sim S}\big[\Delta(x,y)\big] \;=\; 2t - \mathbb{E}_{y\sim S}\Big[\big\langle f(y), \tilde{f} * f(y)\big\rangle_{\mathrm{tr}} + \big\langle \tilde{f} * f(y), f(y)\big\rangle_{\mathrm{tr}}\Big] \tag{6.2}$$

This follows directly by expanding $\Delta(x,y)$ and using the fact that $\|f\|^2 = t$. On the other hand, from the test-passing guarantee we have,

$$\begin{aligned}
\mathop{\mathbb{E}}_{x\sim G, y\sim S}[\Delta(x,y)] \;&\leq\; \Pr_{x\sim G, y\sim S}[\Delta(x,y) > 2\gamma t]\cdot 2t + \Pr_{x\sim G, y\sim S}[\Delta(x,y) \leq 2\gamma t]\cdot 2\gamma t \\
&\leq 2t(1-\delta) + 2\gamma t \\
&=\; 2(1-\delta')t\,.
\end{aligned} \tag{6.3}$$

In the first inequality, we used the fact that $\max_{x,y}\{\Delta(x,y)\} \leq 2t$, and in the second inequality, we used test passing probability to upper bound $\Pr_{x\sim G, y\sim S}[\Delta(x,y) > 2\gamma t]$. Combining Eq. (6.2) and Eq. (6.3), we get:

$$\begin{aligned}
2\delta' t \;&\leq\; \mathbb{E}_{y\sim S}\Big[\big\langle f(y), \tilde{f} * f(y)\big\rangle_{\mathrm{tr}} + \big\langle \tilde{f} * f(y), f(y)\big\rangle_{\mathrm{tr}}\Big] \\
&\leq\; 2\mathop{\mathbb{E}}_{y\sim S}\big[\|f(y)\|_{\mathrm{HS}}\cdot\big\|\tilde{f} * f(y)\big\|_{\mathrm{HS}}\big] && \text{(Cauchy–Schwarz)} \\
&=\; 2\sqrt{t}\cdot\mathop{\mathbb{E}}_{y\sim S}\big[\big\|\tilde{f} * f(y)\big\|_{\mathrm{HS}}\big] && \text{(Using: } f \text{ is unitary-valued.)} \\
&\leq\; 2\sqrt{t}\cdot\Big(\mathop{\mathbb{E}}_{y\sim S}\Big[\big\|\tilde{f} * f(y)\big\|_{\mathrm{HS}}^2\Big]\Big)^{\frac{1}{2}} && \text{(Cauchy–Schwarz)}
\end{aligned}$$

Now we define $\psi(y) = \left\|\tilde{f} * f(y)\right\|_{HS}^2$. Observe that, $\|\psi\|_\infty \leq t$ and it is a positive-definite function by Lemma 6.3.3. This will allow us to deduce that the U^2-norm is large easily. From the computation before, we get,

$$\begin{aligned}
\delta'^2 t &\leq \mathop{\mathbb{E}}_{y\sim S}[\psi(y)] \\
&\leq \mathop{\mathbb{E}}_{y\sim G}[\psi(y)] + \varepsilon\|\psi\|_A && \text{(By Lemma 6.3.1)} \\
&\leq \mathop{\mathbb{E}}_{y\sim G}[\psi(y)] + \varepsilon t && \text{(By Lemma 6.3.3 and Corollary 6.3.2)} \\
&= \mathop{\mathbb{E}}_{y\sim G}\left[\left\|\tilde{f} * f(y)\right\|_{HS}^2\right] + \varepsilon t \\
(\delta'^2 - \varepsilon)\, t &\leq \|f\|_{U^2} && \text{(By definition of } U^2\text{-norm)}
\end{aligned}$$

□

To prove our main result, we need one more component that roughly says that the convolution of functions is mostly supported on low dimensional irreps. The proof is almost identical to the proof of the BNP lemma by Babai, Nikolov, and Pyber [BNP08].

Claim 6.4.3. *Let,* $f, g : G \to M_t(\mathbb{C})$ *and let* $T := \{\rho \in \widehat{G} : d_\rho \geq D\}$*, then the following holds:*

$$\sum_{\rho\in T} d_\rho \left\|\widehat{f * g}\right\|_{HS}^2 \leq \frac{1}{D}\|f\|_2^2\,\|g\|_2^2.$$

In particular, if G *is a* D*-quasirandom group, then* $\left\|f * g - \mu(f * g)\right\| \leq \frac{1}{\sqrt{D}}\|f\|_2\,\|g\|_2$.

Proof. From the convolution identity (Fact 6.2.2), we have:

$$\begin{aligned}
\sum_{\rho\in T} d_\rho \left\|\widehat{f * g}\right\|_{HS}^2 &= \sum_{\rho\in T} d_\rho \left\|\widehat{f}(\rho)\widehat{g}(\rho)\right\|_{HS}^2 && \text{(By convolution identity)} \\
&\leq \sum_{\rho\in T} d_\rho \left\|\widehat{f}(\rho)\right\|_{HS}^2 \left\|\widehat{g}(\rho)\right\|_{HS}^2 && \text{(Norm submultiplicativity)} \\
&\leq \frac{1}{D}\sum_{\rho\in T} d_\rho^2 \left\|\widehat{f}(\rho)\right\|_{HS}^2 \left\|\widehat{g}(\rho)\right\|_{HS}^2 && \text{(Using } d_\rho \geq D \text{ for } \rho \in T\text{)} \\
&\leq \frac{1}{D}\sum_{\rho} d_\rho \left\|\widehat{f}(\rho)\right\|_{HS}^2 \sum_{\rho} d_\rho \left\|\widehat{g}(\rho)\right\|_{HS}^2
\end{aligned}$$

$$= \frac{1}{D}\|f\|_2^2\|g\|_2^2 \qquad \text{(Parseval's identity)}$$

To see the final claim, apply the above to $T := \{\rho \in \widehat{G} : d_\rho \geq D\} = \{\rho \neq \mathrm{triv}\}$ because the group G is D-quasirandom. Using Parseval's identity (Fact 6.2.2) we obtain,

$$\|f * g - \mu(f * g)\|^2 = \sum_{\rho \neq \mathrm{triv}} d_\rho \left\|\widehat{f * g}\right\|_{HS}^2 \leq \frac{1}{D}\|f\|_2^2\|g\|_2^2. \qquad \square$$

Theorem 6.4.4 (Deranomized Homomorphism Testing)**.** *Let G be any finite group and $f : G \to \mathbb{U}_t$ be a unitary matrix-valued function. Let $S \subseteq G$ be an ε-biased set. Assume that f passes the $BLR_{2\gamma}(G, S, f)$ test with probability $\geq \delta(\gamma)$ for any chosen $0 \leq \gamma \leq 1$. Then for $\eta = (\delta - \gamma)^2 - \varepsilon$, the following holds,*

1. *There exists $t' \in [\eta t, \frac{2}{\eta}t]$ and a function $g(x) := V\pi(x)U^*$ where π is a t' dimensional unitary representation, U, V are $t \times t'$ dimensional partial unitary matrices, such that:*
$$\mathbb{E}_x \left\langle f(x), g(x) \right\rangle_{HS} \geq \frac{\eta^2}{4} t$$
2. *for any integer $D > 1$,*
$$\max_{\rho \in \widehat{G} : d_\rho < D} \|\widehat{f}(\rho)\|_{HS}^2 \geq \eta - \frac{t}{D}$$
In particular, there exists an irrep ρ such that $d_\rho < \frac{2t}{\eta}$ such that $\|\widehat{f}(\rho)\|_{HS}^2 \geq \frac{\eta}{2}$.

Proof. From the test passing assumption and Claim 6.4.2, it follows that: $\|f\|_{U^2} \geq \eta t$. Now applying, Theorem 6.4.1 gives us the first claim. For the second claim, our starting point is the same:

$$\eta t \leq \|f\|_{U^2} = \|\tilde{f} * f\|^2 = \sum_{\rho \in \widehat{G}} d_\rho \left\|\widehat{\tilde{f} * f}\,(\rho)\right\|_{HS}^2$$

Now we divide $\widehat{G}$ into low and high dimensional irreps by taking $T := \{\rho \in \widehat{G} : d_\rho \geq D\}$. We have,

$$\eta t \le \sum_{\rho\in\widehat{G}} d_\rho \left\|\widehat{\tilde{f} * f(\rho)}\right\|_{HS}^2 = \sum_{\rho\in\widehat{G}\setminus T} d_\rho \left\|\widehat{\tilde{f} * f}(\rho)\right\|_{HS}^2 + \sum_{\rho\in T} d_\rho \left\|\widehat{\tilde{f} * f}(\rho)\right\|_{HS}^2$$
$$\le \sum_{\rho\in\widehat{G}\setminus T} d_\rho \left\|\widehat{\tilde{f} * f}(\rho)\right\|_{HS}^2 + \frac{\|\tilde{f}\|^2\|f\|^2}{D}.$$

In the inequality step above, we used Claim 6.4.3. As f is unitary valued, $\|f\|^2 = \|\tilde{f}\|^2 = t$. It follows that: $\sum_{\rho\in\widehat{G}\setminus T} d_\rho \left\|\widehat{\tilde{f} * f}(\rho)\right\|_{HS}^2 \ge \eta t - t^2/D$. Finally, we have,

$$\begin{aligned}
\eta t - t^2/D &\le \sum_{\rho\in\widehat{G}\setminus T} d_\rho \left\|\widehat{\tilde{f} * f}(\rho)\right\|_{HS}^2 \\
&= \sum_{\rho\in\widehat{G}\setminus T} d_\rho \left\|\widehat{f}(\rho)^* \widehat{f}(\rho)\right\|_{HS}^2 && \text{(Convolution identity, Fact 6.2.2)} \\
&\le \sum_{\rho\in\widehat{G}\setminus T} d_\rho \|\widehat{f}(\rho)\|_{HS}^4 && \text{(Sub-Multiplicativity.)} \\
&\le \max_{\rho\in\widehat{G}: d_\rho < D} \|\widehat{f}(\rho)\|_{HS}^2 \cdot \sum_{\rho} d_\rho \|\widehat{f}(\rho)\|_{HS}^2 && \text{(As, } d_\rho \le D \text{ holds for any } \rho \in \widehat{G}\setminus T) \\
&= \max_{\rho\in\widehat{G}: d_\rho < D} \|\widehat{f}(\rho)\|_{HS}^2 \cdot \|f\|^2 && \text{(Parseval's identity, Fact 6.2.2)} \\
&= t \cdot \max_{\rho\in\widehat{G}: d_\rho < D} \|\widehat{f}(\rho)\|_{HS}^2 && (f \text{ is unitary, } \|f\|^2 = t). \qquad \square
\end{aligned}$$

6.5 Derandomized Mixing

In this section, we prove a general "degree–2 mixing lemma" as explained in the introduction for the general case of matrix-valued functions. The assumption that the functions are mean-zero is without loss of generality and only for brevity.

Lemma 6.1.6 (Derandomized Matrix BNP)**.** *Let G be a group such that the dimension of the smallest non-trivial irrep is D, and let $S \subseteq G$ be an ε-biased set. Let $f, g : G \to M_t(\mathbb{C})$ be*

mean-zero functions. Then,

$$\mathop{\mathbb{E}}_{s\sim S}\left[\,\|(f*g)(s)\|_{\mathrm{HS}}^2\right] \;\le\; \left(\frac{1}{D}+\varepsilon\right)\|f\|_2^2\|g\|_2^2$$

The usual BNP lemma can be recovered by setting $S = G$*, and thus,* $\varepsilon = 0$*.*

To prove this derandomization, we first prove a more general version of Lemma 6.3.3, where instead of $\psi(y) = \left\|(\tilde{f}*f)(y)\right\|_{\mathrm{HS}}^2$, we have $\psi(y) = \|(f*g)(y)\|_{\mathrm{HS}}^2$. This is no longer positive definite as earlier. Still, we can use elementary representation theory to explicitly give a factorization of the form $\psi(y) = \langle u, \phi(y)\, v\rangle$ and thereby compute the algebra norm. The representation ϕ that will come up is defined as follows — let $V \subseteq \mathcal{L}_t^2(G\times G)$ be the subspace $V = \mathrm{span}\left\{F(x,y) = f(y)f(x)^* \mid f \in \mathcal{L}_t^2(G)\right\}$. Note that V inherits the expectation trace inner product, i.e., $\langle F, H\rangle = \mathbb{E}_{x,y\sim G}[\langle F(x,y), H(x,y)\rangle_{\mathrm{tr}}]$. Then, we can define the following representation of G,

$$(\phi(a)\cdot F)(x,y) = F(xa,ya) = f(ya)f(xa)^*.$$

We are ready to prove the generalization of Lemma 6.3.3.

Lemma 6.5.1. *Let* $f, h : G \to M_t(\mathbb{C})$ *be matrix-valued functions, and define* $\psi(y) = \|(h*f)(y)\|_{\mathrm{HS}}^2$*. Then,* $\|\psi\|_A \le \|f\|^2\|h\|^2$*. Moreover, if the functions are unitary-valued,* $\|\psi\|_A \le t$*.*

Proof. Let $F(x,y) = f(y)f(x)^*$ and similarly, $\tilde{H}(x,y) = \tilde{h}(y)\tilde{h}(x)^*$. Now,

$$\begin{aligned}
\psi(a) &= \left\|\mathop{\mathbb{E}}_{x\sim G}\left[h(x^{-1})f(xa)\right]\right\|^2 \\
&= \left\langle \mathop{\mathbb{E}}_{x\sim G}\left[\tilde{h}(x)^*f(xa)\right], \mathop{\mathbb{E}}_{y\sim G}\left[\tilde{h}(y)^*f(ya)\right]\right\rangle_{\mathrm{tr}} \\
&= \mathop{\mathbb{E}}_{x,y\sim G}\left[\left\langle \tilde{h}(x)^*f(xa), \tilde{h}(y)^*f(ya)\right\rangle_{\mathrm{tr}}\right] \\
&= \mathop{\mathbb{E}}_{x,y\sim G}\left[\left\langle \tilde{h}(y)\tilde{h}(x)^*, f(ya)f(xa)^*\right\rangle_{\mathrm{tr}}\right] \\
&= \mathop{\mathbb{E}}_{x,y\sim G}\left[\left\langle \widetilde{H}(x,y), \phi(a)\,F(x,y)\right\rangle_{\mathrm{tr}}\right]
\end{aligned}$$

$$= \left\langle \widetilde{H}, \phi(a)\, F \right\rangle.$$

By the definition of algebra norm Definition 6.2.3, we have $\|\psi\|_A \leq \|H\| \cdot \|F\|$. Now,

$$\|F\|^2 = \mathop{\mathbb{E}}_{x,y}\left[\|f(y)f(x)^*\|_{HS}^2\right] \leq \mathop{\mathbb{E}}_{x,y}\left[\|f(y)\|_{HS}^2 \|f(x)^*\|_{HS}^2\right] = \|f\|^4.$$

The inequality here follows from sub-multiplicativity. Similarly, $\|\widetilde{H}\| = \|h\|^2$. This proves the first claim. When the functions map to unitary matrices, $\widetilde{H}(x,y), F(x,y)$ are unitary and thus, $\left\|\widetilde{H}(x,y)\right\|_{HS}^2 = \|F(x,y)\|_{HS}^2 = t$ for every $x, y \in G$. We can thus avoid using sub-multiplicativity and directly obtain,

$$\|F\|^2 = \mathop{\mathbb{E}}_{x,y}\left[\|f(y)f(x)^*\|_{HS}^2\right] = t = \|\widetilde{H}\|^2. \qquad \square$$

Remark 6.5.2. One can weaken the unitary assumption by requiring that f is "unitary on average", an assumption also used in [MR15]. This is because $\|F\| = \|\mathbb{E}_x[f(x)f(x)^*]\|_{HS}$.

Lemma 6.1.6 (Derandomized Matrix BNP)**.** *Let G be a group such that the dimension of the smallest non-trivial irrep is D, and let $S \subseteq G$ be an ε-biased set. Let $f, g : G \to M_t(\mathbb{C})$ be mean-zero functions. Then,*

$$\mathop{\mathbb{E}}_{s\sim S}\left[\, \|(f * g)(s)\|_{HS}^2\right] \;\leq\; \left(\frac{1}{D} + \varepsilon\right) \|f\|_2^2 \|g\|_2^2$$

The usual BNP lemma can be recovered by setting $S = G$, and thus, $\varepsilon = 0$.

Proof. From Lemma 6.5.1, we have that the function, $\psi(s) := \|(f * g)(s)\|_{HS}^2$ has algebra norm $\|f\|_2^2 \|g\|_2^2$ and therefore from Lemma 6.3.1 we have that averaging over S is ε-close to true average. Thus,

$$\begin{aligned}
\mathop{\mathbb{E}}_{s\sim S}\left[\|(f * g)(s)\|_{HS}^2\right] &\leq \mathop{\mathbb{E}}_{s\sim G}\left[\|(f * g)(s)\|_{HS}^2\right] + \varepsilon\|f\|_2^2 \|g\|_2^2 && \text{(Using Lemma 6.3.1)} \\
&\leq \frac{1}{D}\|f\|_2^2 \|g\|_2^2 + \varepsilon\|f\|_2^2 \|g\|_2^2 && \text{(Using Claim 6.4.3)} \qquad \square
\end{aligned}$$

CHAPTER 7

CONCLUSION

This thesis has used iterative techniques—graph lifts and derandomized powering—to build expander graphs with symmetries of various families of groups. We have also seen how such graphs can be used to build quantum error correction codes and randomness-efficient testing algorithms. Apart from natural questions regarding the quantitative strengthening of our results, there are several exciting avenues for further exploration.

High-Dimensional Expanders The notion of expansion has been generalized to hypergraphs, and just like an expander is a well-connected graph, an *high-dimensional expander* (HDX) is a well-connected sparse hypergraph. Hypergraphs have been well-studied, but the concept of an HDX is a recent development, and much needs to be learned about these objects. A different perspective is that HDXs are graphs with an additional inductive structure that allows for *local-to-global* theorems. For instance, an HDX is an expander if all its "local subgraphs" are. Such a property is very useful, and HDXs have been instrumental in proving mixing times of well-known Markov chains [ALGV24], solving constraint satisfaction problems [AJT19], and in derandomized *agreement testing* [DDL24, BLM24]. However, the iterative techniques used to build graphs do not easily generalize to HDXs. There have been some works on generalizing *graph powering* [KP21] and *graph lifts* [BYDM24] to HDXs. Still, much needs to be known, and designing expander construction techniques that preserve this simplicial structure is a very interesting direction.

Non-Abelian Symmetries For most families of groups, the best construction of Cayley expanders we have is of logarithmic degree from Alon and Roichman [AR94] and its derandomizations. A central open question is to determine which groups admit constant-degree expanding Cayley graphs and give efficient constructions if they do. One important

family that admits these is *non-abelian finite simple groups*, and even constantly many random generators yield an expander. However, the proof of this using the Bourgain–Gamburd machinery [BG08] is quite involved [BGGT15], and one could hope for an elementary proof using recent advances in random matrix theory [BBvH23]. A *square Cayley complex* is the central object underlying the breakthrough constructions of quantum LDPC codes and locally testable codes. The underlying graph is a product of Cayley graphs with a pair of commuting actions over non-abelian groups. Dinur, Lin, and Vidick [DLV24] obtained almost good quantum locally testable codes via a construction with a larger number of commuting actions of Ablelian groups. Having such a construction via non-Abelian groups might yield good quantum LTCs.

Study of functions on groups The study of boolean functions, i.e., $f : \mathbb{Z}_2^n \to \mathbb{R}$, has had applications in multiple areas like complexity theory, learning theory, etc (see [O'D14]). It is fruitful, however, to go beyond Boolean functions and generalize the machinery of boolean function analysis to matrix-valued functions over general groups, $f : G \to \mathbb{C}^{n\times n}$. For instance, in Chapter 6, we generalized the notion of *spectral norm* and its connection to ε-biased sets. A few interesting lines of inquiry based on recent applications are,

- Generalizing *hypercontractive inequalities* to non-abelian groups such as compact groups and the symmetric group, has applications in quantum complexity [AGL23], in extremal combinatorics [FKLM24], and potentially in many other settings.

- Abelian groups are poorly mixing, and thus, the question of *mixing in groups* is exciting in the non-abelian world [BT14, Tao13]. Such mixing results are also studied in additive combinatorics due to connections to the density of arithmetic progressions and are also a crucial ingredient in recent results on optimal hardness of approximability [BK21].

REFERENCES

[ACKM19] Naman Agarwal, Karthekeyan Chandrasekaran, Alexandra Kolla, and Vivek Madan. On the Expansion of Group-Based Lifts. *SIAM J. Discret. Math.*, 33(3):1338–1373, 2019.

[AGHP92] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

[AGL23] Srinivasan Arunachalam, Uma Girish, and Noam Lifshitz. One Clean Qubit Suffices for Quantum Communication Advantage, 2023. Preprint, `2310.02406 [quant-ph]`.

[AHL02] Noga Alon, Shlomo Hoory, and Nathan Linial. The Moore bound for irregular graphs. *Graphs and Combinatorics*, 18:53–57, 2002.

[AHL+14] Dorit Aharonov, Aram W. Harrow, Zeph Landau, Daniel Nagaj, Mario Szegedy, and Umesh V. Vazirani. Local Tests of Global Entanglement and a Counterexample to the Generalized Area Law. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science*, 2014.

[AJT19] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani. Approximating constraint satisfaction problems on high-dimensional expanders. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science*, pages 180–201, 2019.

[AL02] Alon Amit and Nathan Linial. Random Graph Coverings I: General Theory and Graph Connectivity. *Combinatorica*, 22(1):1–18, 2002.

[AL06] Alon Amit and Nathan Linial. Random Lifts of Graphs: Edge Expansion. *Combinatorics, Probability and Computing*, 15(3):317–332, 2006.

[ALGV24] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. *Annals of Mathematics*, 199(1), 2024.

[ALMR01] Alon Amit, Nathan Linial, Jiří Matoušek, and Eyal Rozenman. Random lifts of graphs. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 883–894, 2001.

[Alo21] Noga Alon. Explicit Expanders of Every Degree and Size. *Combinatorica*, 41:447–463, 2021.

[AR94] Noga Alon and Yuval Roichman. Random cayley graphs and expanders. *Random Struct. Algorithms*, 5(2):271–285, 1994.

[AS04] Andris Ambainis and Adam D. Smith. Small pseudo-random families of matrices: Derandomizing approximate quantum encryption. In *APPROX-RANDOM 2004 Proceedings*, volume 3122 of *Lecture Notes in Computer Science*, pages 249–260. Springer, 2004.

[BASTS08] Avraham Ben-Aroya, Oded Schwartz, and Amnon Ta-Shma. Quantum expanders: Motivation and constructions. In *Proceedings of the 23rd IEEE Conference on Computational Complexity*, pages 292–303, 2008.

[BATS08] Avraham Ben-Aroya and Amnon Ta-Shma. A Combinatorial Construction of Almost-Ramanujan Graphs using the Zig-Zag Product. In *Proceedings of the 40th ACM Symposium on Theory of Computing*, pages 325–334, 2008.

[BBvH23] Afonso S. Bandeira, March T. Boedihardjo, and Ramon van Handel. Matrix concentration inequalities and free probability. *Inventiones mathematicae*, 234(1):419–487, 2023.

[BCH+95] M. Bellare, D. Coppersmith, J. Håstad, M. Kiwi, and M. Sudan. Linearity testing in characteristic two. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 432–441, 1995.

[Bea97] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 48–53, 1997.

[BFL03] László Babai, Katalin Friedl, and András Lukács. Near representations of finite groups, 2003. Manuscript.

[BG08] Jean Bourgain and Alex Gamburd. Uniform expansion bounds for Cayley graphs of $\mathrm{SL}_2(\mathbb{F}_p)$. *Annals of Mathematics*, 167(2):625–642, 2008.

[BGGT15] Emmanuel Breuillard, Ben J. Green, Robert M. Guralnick, and Terence Tao. Expansion in finite simple groups of Lie type. *Journal of the European Mathematical Society*, 17(6):1367–1434, 2015.

[BHR22] Amey Bhangale, Prahladh Harsha, and Sourya Roy. Mixing of 3-term progressions in Quasirandom Groups. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference*, pages 20:1–20:9, 2022.

[BISW01] B. Barak, R. Impagliazzo, A. Shpilka, and A. Wigderson. Dimension Expanders. Manuscript, 2001.

[BK21] Amey Bhangale and Subhash Khot. Optimal inapproximability of satisfiable k-LIN over non-abelian groups. In *Proceedings of the 53rd ACM Symposium on Theory of Computing*, 2021.

[BKL89] László Babai, William M. Kantor, and A. Lubotzky. Small-diameter Cayley Graphs for Finite Simple Groups. *Eur. J. Comb.*, 10, 1989.

[BKM22] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k-CSPs: I. In *Proceedings of the 54th ACM Symposium on Theory of Computing*, 2022.

[BL06] Yonatan Bilu and Nathan Linial. Lifts, Discrepancy and Nearly Optimal Spectral Gap. *Combinatorica*, 26(5):495–519, 2006.

[BL22] Emmanuel Breuillard and Alexander Lubotzky. Expansion in simple groups. In *Dynamics, Geometry, Number Theory*, pages 246–275, 2022.

[BLM24] Mitali Bafna, Noam Lifshitz, and Dor Minzer. Constant Degree Direct Product Testers with Small Soundness. In *Proceedings of the 65h IEEE Symposium on Foundations of Computer Science*, 2024.

[BLR90] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. In *Proceedings of the 22nd ACM Symposium on Theory of Computing*, pages 73–83, 1990.

[BM06] L.M.J. Bazzi and S.K. Mitter. Some randomized code constructions from group actions. *IEEE Transactions on Information Theory*, 52(7):3210–3219, 2006.

[BNP08] László Babai, Nikolay Nikolov, and László Pyber. Product growth and mixing in finite groups. In *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms*, 2008.

[BOCLR07] Michael Ben-Or, Don Coppersmith, Mike Luby, and Ronitt Rubinfeld. Non-Abelian homomorphism testing, and distributions close to their self-convolutions. *Random Structures & Algorithms*, 32(1):49–70, 2007.

[Bor20] Charles Bordenave. A new proof of Friedman's second eigenvalue theorem and its extension to random lifts. *Annales scientifiques de lÉcole normale supérieure*, 53(6):1393–1439, 2020.

[BP18] Roman Badora and Barbara Przebieracz. On approximate group homomorphisms. *Journal of Mathematical Analysis and Applications*, 2018.

[BS88] László Babai and Akos Seress. On the Diameter of Cayley Graphs of the Symmetric Group. *Journal of Combinatorial Theory, Series A*, 49(1), 1988.

[BSS05] László Babai, Amir Shpilka, and Daniel Stefankovic. Locally testable cyclic codes. *IEEE Trans. Inf. Theory*, 51(8):2849–2858, 2005.

[BSVW03] Eli Ben-Sasson, Madhu Sudan, Salil P. Vadhan, and Avi Wigderson. Randomness-efficient low degree tests and short PCPs via ε-biased sets. In *Proceedings of the 35th ACM Symposium on Theory of Computing*, pages 612–621, 2003.

[BT14] Vitaly Bergelson and Terence Tao. Multiple Recurrence in Quasirandom Groups. *Geometric and Functional Analysis*, 24(1):1–48, 2014.

[BY13] Jean Bourgain and Amir Yehudayoff. Expansion in $\mathrm{SL}_2(\mathbb{R})$ and monotone expanders. *Geometric and Functional Analysis*, 23(1), 2013.

[BYDM24] Inbar Ben Yaacov, Yotam Dikstein, and Gal Maor. Sparse High Dimensional Expanders via Local Lifts, 2024. Preprint, `arXiv:2405.19191 [cs.DM]`.

[Che10] Yuan-You Fu-Rui Cheng. Explicit estimate on primes between consecutive cubes. *Rocky Mountain Journal of Mathematics*, 40(1), 2010.

[CMR13] Sixia Chen, Cristopher Moore, and Alexander Russell. Small-bias sets for nonabelian groups - derandomizations of the Alon–Roichman theorem. In *APPROX-RANDOM*, volume 8096 of *Lecture Notes in Computer Science*, pages 436–451, 2013.

[Con] Keith Conrad. Simultaneous Ccommutativity Of Operators. `https://kconrad.math.uconn.edu/blurbs/linmultialg/simulcomm.pdf`. [Online; accessed 7-September-2021].

[CPW69] C.L. Chen, W.W. Peterson, and E.J. Weldon. Some results on quasi-cyclic codes. *Information and Control*, 15(5):407–423, 1969.

[CS96] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Phys. Rev. A*, 54:1098–1105, 1996.

[DCOT18] Marcus De Chiffre, Narutaka Ozawa, and Andreas Thom. Operator algebraic approach to inverse and stability theorems for amenable groups. *Mathematika*, 65(1):98–118, 2018.

[DDL24] Yotam Dikstein, Irit Dinur, and Alexander Lubotzky. Low Acceptance Agreement Tests via Bounded-Degree Symplectic HDXs, 2024. Preprint, `arXiv:2402.01078 [cs.CCs]`.

[DETT10] Anindya De, Omid Etesami, Luca Trevisan, and Madhur Tulsiani. Improved Pseudorandom generators for depth 2 circuits. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2010.

[DHLV23] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good Quantum LDPC Codes with Linear Time Decoders. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. ACM, 2023.

[DLV24] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of higher-dimensional cubical complexes with application to quantum locally testable codes. 2024.

[DS09] Zeev Dvir and Amir Shpilka. Towards dimension expanders over finite fields. *Combinatorica*, 31(3), 2009.

[DW10] Zeev Dvir and Avi Wigderson. Monotone expanders: Constructions and applications. *Theory of Computing*, 6(12), 2010.

[EKZ20] Shai Evra, Tali Kaufman, and Gilles Zémor. Decodable quantum LDPC codes beyond the $\sqrt{n}$ distance barrier using high dimensional expanders. In *Proceedings of the 61st IEEE Symposium on Foundations of Computer Science*, pages 218–227, 2020.

[Eym64] P. Eymard. L'algébre de Fourier d'un groupe localement compact. *Bulletin de la Société Mathématique de France*, 92:181–236, 1964.

[Far00] Ilijas Farah. Approximate homomorphisms II: Group homomorphisms. *Combinatorica*, 20(1):47–60, 2000.

[FG15] Michael A. Forbes and Venkatesan Guruswami. Dimension Expanders via Rank Condensers. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, volume 40, pages 800–814, 2015.

[FKLM24] Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer. Hypercontractivity on the symmetric group. *Forum of Mathematics, Sigma*, 12, 2024.

[Fri03] Joel Friedman. A proof of Alon's second eigenvalue conjecture. In *Proceedings of the 35th ACM Symposium on Theory of Computing*, 2003.

[GGMT23] W. T. Gowers, Ben Green, Freddie Manners, and Terence Tao. On a conjecture of Marton, 2023. Preprint, `arXiv:2311.05762 [math.NT]`.

[GH17] William Timothy Gowers and Omid Hatami. Inverse and stability theorems for approximate representations of finite groups. *Sbornik: Mathematics*, 208(12):1784, 2017.

[Gow08] W. T. Gowers. Quasirandom Groups. *Combinatorics, Probability and Computing*, 17(3):363–387, 2008.

[GS08] Benjamin Green and Tom Sanders. A quantitative version of the idempotent theorem in harmonic analysis. *Annals of Mathematics*, 168(3):1025–1054, 2008.

[Har07] Aram W. Harrow. Quantum expanders from any classical Cayley graph expander. *Quantum Information & Computation*, 2007.

[Has07a] M. B. Hastings. Entropy and entanglement in quantum ground states. *Physical Review B*, 76(3), 2007.

[Has07b] M. B. Hastings. Random unitaries give quantum expanders. *Phys. Rev. A*, 76:032315, 2007.

[HH09] M. B. Hastings and A. W. Harrow. Classical and quantum tensor product expanders. *Quantum Info. Comput.*, 2009.

[HHH22] Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. Dimension-free bounds and structural results in communication complexity. *Israel Journal of Mathematics*, 253(2):555–616, 2022.

[HHO21] Matthew B. Hastings, Jeongwan Haah, and Ryan O'Donnell. Fiber bundle codes: breaking the $n^{1/2}\mathrm{polylog}(n)$ barrier for quantum LDPC codes. In *Proceedings of the 53rd ACM Symposium on Theory of Computing*, pages 1276–1288. ACM, 2021.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43(04):439–562, 2006.

[HPS16] Chris Hall, Doron Puder, and William F. Sawin. Ramanujan coverings of graphs. In *Proceedings of the 48th ACM Symposium on Theory of Computing*, pages 533–541, 2016.

[HW03] Johan Håstad and Avi Wigderson. Simple analysis of graph tests for linearity and PCP. *Random Structures & Algorithms*, 22(2):139–160, 2003.

[JM21] Akhil Jalan and Dana Moshkovitz. Near-Optimal Cayley Expanders for Abelian Groups. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021)*, volume 213, pages 24:1–24:23, 2021.

[JNV+21] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP* = RE. *Commun. ACM*, 64(11):131–138, 2021.

[JQST20] Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. Unique decoding of explicit ε-balanced codes near the Gilbert–Varshamov bound. In *Proceedings of the 61st IEEE Symposium on Foundations of Computer Science*, 2020.

[Kas07] Martin Kassabov. Symmetric groups and expander graphs. *Inventiones mathematicae*, 170(2):327–354, 2007.

[Kit97] A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, 1997.

[Kiw03] M. Kiwi. Algebraic testing and weight distributions of codes. *Theoretical Computer Science*, 299(1):81–106, 2003.

[KKN21] Marek Kaluba, Dawid Kielak, and Piotr W. Nowak. On property (T) for $\mathrm{Aut}(\mathbb{F}_n)$ and $\mathrm{SL}_n(\mathbb{Z})$. *Annals of Mathematics*, 193(2):539 – 562, 2021.

[KM93] Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the fourier spectrum. *SIAM Journal on Computing*, 22(6):1331–1348, 1993.

[KM23] Z. Kelley and R. Meka. Strong bounds for 3-progressions. In *Proceedings of the 64th IEEE Symposium on Foundations of Computer Science*, pages 933–973, 2023.

[KN06] Martin Kassabov and Nikolay Nikolov. Universal lattices and Property τ. *Inventiones mathematicae*, 165(1):209–224, 2006.

[KP21] Tali Kaufman and Ori Parzanchevski. Free Flags over Local Rings and Powering of High Dimensional Expanders. *International Mathematics Research Notices*, 2021.

[KT21] Tali Kaufman and Ran J. Tessler. New cosystolic expanders from tensors imply explicit quantum LDPC codes with $\Omega(\sqrt{n}\log^k n)$ distance. In *Proceedings of the 53rd ACM Symposium on Theory of Computing*, pages 1317–1329. ACM, 2021.

[LBM+18] Huaan Li, Baoming Bai, Xijin Mu, Ji Zhang, and Hengzhou Xu. Algebra-assisted construction of quasi-cyclic LDPC codes for 5G new radio. *IEEE Access*, 6:50229–50244, 2018.

[LP00] Alexander Lubotzky and Igor Pak. The product replacement algorithm and Kazhdan's property (T). *Journal of the American Mathematical Society*, 14(2):347–363, 2000.

[LPS88] Alexander Lubotzky, R. Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8:261–277, 1988.

[Lub11] Alexander Lubotzky. Finite simple groups of Lie type as expanders. *Journal of the European Mathematical Society*, pages 1331–1341, 2011.

[Lub12] Alexander Lubotzky. Expander graphs in pure and applied mathematics. *Bull. Amer. Math. Soc.*, 49:113–162, 2012.

[LZ08] Alexander Lubotzky and Efim Zelmanov. Dimension expanders. *Journal of Algebra*, 319(2):730–738, 2008.

[LZ22] Anthony Leverrier and Gilles Zémor. Quantum Tanner codes. In *Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science*, 2022.

[Mar88] G. A. Margulis. Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problemy Peredachi Informatsii*, 24(1):51–60, 1988.

[MOP20] Sidhanth Mohanty, Ryan O'Donnell, and Pedro Paredes. Explicit near-Ramanujan graphs of every degree. In *Proceedings of the 52nd ACM Symposium on Theory of Computing*, pages 510–523. ACM, 2020.

[Mor94] M. Morgenstern. Existence and explicit constructions of q + 1 regular Ramanujan graphs for every prime power q. *J. Comb. Theory Ser. B*, pages 44–62, 1994.

[MR15] Cristopher Moore and Alexander Russell. Approximate representations, approximate homomorphisms, and low-dimensional embeddings of groups. *SIAM Journal on Discrete Mathematics*, 29(1):182–197, 2015.

[MSS15] Adam Marcus, Daniel Spielman, and Nikhil Srivastava. Interlacing families i: Bipartite Ramanujan graphs of all degrees. *Annals of Mathematics*, 2015.

[Nil91] Alon Nilli. On the second eigenvalue of a graph. *Discrete Mathematics*, 91(2):207–210, 1991.

[NN90] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. In *Proceedings of the 22nd ACM Symposium on Theory of Computing*, pages 213–223, 1990.

[NN93] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.

[NV17] Anand Natarajan and Thomas Vidick. A quantum linearity test for robustly verifying entanglement. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, 2017.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 1 edition, 2014.

[OW20] Ryan O'Donnell and Xinyu Wu. Explicit near-fully X-Ramanujan graphs. In *Proceedings of the 61st IEEE Symposium on Foundations of Computer Science*, pages 1045–1056, 2020.

[PK21] Pavel Panteleev and Gleb Kalachev. Quantum LDPC Codes with Almost Linear Minimum Distance. *IEEE Transactions on Information Theory*, 2021.

[PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *Proceedings of the 54th ACM Symposium on Theory of Computing*, pages 375–388, 2022.

[PZ02] Igor Pak and Andrzej Zuk. On Kazhdan constants and Mixing of Random Walks. *International Mathematics Research Notices*, 36:1891–1905, 2002.

[Rao19] Shravas Rao. A Hoeffding inequality for Markov chains. *Electronic Communications in Probability*, 24:1 – 11, 2019.

[Rei05] Omer Reingold. Undirected ST-connectivity in log-space. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 376–385, 2005.

[RL10] J.D. Rogawski and A. Lubotzky. *Discrete Groups, Expanding Graphs and Invariant Measures*. Modern Birkhäuser Classics. Birkhäuser Basel, 2010.

[RVW00] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, 2000.

[Sam07] A. Samorodnitsky. Low-degree tests at large distances. In *Proceedings of the 39th ACM Symposium on Theory of Computing*, pages 506–515, 2007.

[San11] Tom Sanders. A Quantitative Version of the Non-Abelian Idempotent Theorem. *Geometric and Functional Analysis*, 21(1):141–221, 2011.

[San12] Tom Sanders. On the Bogolyubov-Ruzsa lemma. *Analysis & PDE*, 5(3):627–655, 2012.

[San21] Tom Sanders. Coset decision trees and the Fourier algebra. *Journal d'Analyse Mathématique*, 144(1):227–259, 2021.

[ST00] Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.

[Ste96] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 452(1954):2551–2577, 1996.

[STV17] Amir Shpilka, Avishay Tal, and Ben Lee Volk. On the Structure of Boolean functions with Small Spectral Norm. *Computational Complexity*, 26(1):229–273, 2017.

[SW04] Amir Shpilka and Avi Wigderson. Derandomizing Homomorphism Testing in General Groups. In *Proceedings of the 36th ACM Symposium on Theory of Computing*, 2004.

[Tao13] Terence Tao. Mixing for progressions in non-abelian groups. *Forum of Mathematics, Sigma*, 1, 2013.

[TS17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, pages 238–251, 2017.

[TZ14] Jean-Pierre Tillich and Gilles Zémor. Quantum LDPC codes with positive rate and minimum distance proportional to $n^{1/2}$. *IEEE Transactions on Information Theory*, 60(2):1193–1202, 2014.

[Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2012.

[WX08] Avi Wigderson and David Xiao. Derandomizing the Ahlswede-Winter matrix-valued Chernoff bound using pessimistic estimators, and applications. *Theory Comput.*, 4(1):53–76, 2008.

APPENDIX A

APPENDIX

A.1 Instantiating the s-wide Replacement Product

The goal of this section is to prove the following result that implies Theorem 1.2.2.

Theorem A.1.1 (Almost Ramanujan Expanders I)**.** *Let* $\mathrm{Cay}(G, S)$ *be* λ_0*-expander with constant* $\lambda_0 \in (0, 1)$*. For every function* $\beta(\lambda) > 0$*, and for any* $\lambda > 0$*, sufficiently small such that*

$$\frac{32}{\beta(\lambda)} \leq \left(\frac{\log(1/\lambda)}{4\log\log(1/\lambda)}\right)^{1/3},$$

there exists a deterministic polynomial time algorithm to construct S' *such that* $\mathrm{Cay}(G, S')$ *is a* λ*-expander with degree* $|S'| = O_{\lambda_0}(|S|/\lambda^{2+\beta})$*.*

Furthermore, each element in S' *is the product of* $O(\log(1/\lambda))$ *elements of* S*.*

Overview

We will explicitly construct the graphs X and Y as needed for the s-wide product. Once we obtain the graphs, we identify the vertices of X, i.e., V_X with the initial generating set S, or perhaps a slightly modified set S', obtained by duplicating and adding identities. The final set is obtained by multiplying elements along each $(t-1)$-length walk on the s-wide replacement product of X and Y. The way we choose parameters and objects for it borrows heavily from Ta-Shma's arguments in [TS17]. The analysis follows an analogous structure of [JQST20] for binary codes, which, in turn, builds on the original analysis of Ta-Shma [TS17]. We will also use the following result from that work,

Lemma A.1.2 (Based on Lemma 6 [TS17])**.** *For every* $m \in \mathbb{N}^+$ *and* $d = 2^{2k} \leq 2^m$*, there exists a fully explicit set* $A \subseteq \mathbb{Z}_2^m$ *such that the graph* $\mathrm{Cay}(\mathbb{Z}_2^m, A)$ *is a* $(2^m, d, \lambda = \frac{m}{\sqrt{d}})$*-expander graph.*

The construction Given as input $n := |S|, \lambda$ and a slowly growing function $\beta(\lambda)$, we construct the graphs X, Y as described below with the parameters $(s, d_1, d_2, \lambda_1, \lambda_2)$ which are all functions of λ and $\beta(\lambda)$. These are summarized in a table below for reference[1]. Recall that a (n, d, λ)-graph has n vertices, is d-regular, and has the second largest singular value of its normalized adjacency matrix at most λ.

- **Outer Graph** — The outer graph X will be an (n', d_1, λ_1)-graph which is a Cayley graph on $SL_2(p)$ constructed using Corollary 4.2.11 with (n, λ_1) as input. By Example 4.3.4, we obtain a locally invertible graph on $n' \approx n$. The condition on the size is satisfied as $n = 2|S|d_2^5 \geq d_2^5 \geq 2^{2^{17}}$ by the assumption that $s \geq 2^{10}$. Moreover, the degree is $\frac{c}{\lambda_1^{2*4.1}} \leq \frac{c d_2^{4.1}}{b_2^{8.2}} \leq d_2^5$. We increase its degree to d_2^5 by taking multiple copies of the generating set, which does not change bias[2]. Thus, we obtain a (n', d_1, λ_1)-graph where $n' = n + O(n^{8/9})$.

- **Inner Graph** — The inner graph Y will be a (d_1^s, d_2, λ_2)-graph which is a Cayley graph on $\mathbb{Z}_2^m$ and therefore by Lemma 4.3.9, it is compatible. For this, we use the construction of Alon et al. [AGHP92] and the analysis of Ta-Shma (Lemma A.1.2).

We summarize the construction and the choice of parameters $n', d_1, d_2, \lambda_1, \lambda_2$ and s, which are chosen as follows for a fixed $\beta(\lambda)$ here -

s is the smallest power of 2 such that $\frac{32}{\beta} \leq s \leq \left(\frac{\log(1/\lambda)}{4 \log\log(1/\lambda)}\right)^{1/3}$
Every other parameter is a function of s.
$Y : (n_2, d_2, \lambda_2), \quad n_2 = d_2^{5s}, \quad d_2 = s^{4s}, \quad \lambda_2 \leq \frac{b_2}{\sqrt{d_2}}, \quad b_2 = 5s \log d_2$
$X : (n', d_1, \lambda_1), \quad n' \approx n = O(\lvert S\rvert d_2^5), \quad d_1 = d_2^5, \quad \lambda_1 = \frac{\lambda_2^2}{10}$
$t :$ smallest integer such that $(\lambda_2)^{(1-5\alpha)(1-\alpha)(t-1)} \leq \lambda, ;$ where $\alpha = 1/s$

1. The choice of parameters is similar but not identical to Ta-Shma's choice.

2. This is wasteful, but we do it to ensure that $V(Y) = d_1^s$ and that d_1^s is a power of 2.

Note: We assume that $s \geq 2^{10}$ since otherwise λ is a constant and we can use Theorem 4.2.2. Now, we mention the central claim that we need from our choice of parameters.

Claim A.1.3. *The selection of the parameters above implies the following bounds on* t*,*

(i). $t - 1 \geq 2s^2$

(ii). $(d_2)^{(t-1)} \leq \lambda^{-2(1+10\alpha)}$,

Proof. Proof of (i) : Using $d_2 = s^{4s}$ and the upper bound on s, we have

$$\left(\frac{1}{\lambda_2}\right)^{(1-5\alpha)(1-\alpha)2s^2} \leq \left(\frac{1}{\lambda_2}\right)^{2s^2} = \left(\frac{d_2}{b_2^2}\right)^{s^2} \leq (d_2)^{s^2} = s^{4s^3}$$
$$= 2^{4s^3 \log_2(s)} \leq 2^{\log_2(1/\lambda)} = \frac{1}{\lambda}.$$

Hence, $(\lambda_2)^{(1-5\alpha)(1-\alpha)s/\alpha} \geq \lambda$ and thus $t - 1$ must be at least $2s^2$. Also, observe that,

$$\lambda_2^{(1-5\alpha)(1-\alpha)^2(t-1)} = \lambda_2^{(1-5\alpha)(1-\alpha)(t-2)\left(\frac{(1-\alpha)}{1-1/(t-1)}\right)} \tag{A.1}$$
$$\geq \lambda_2^{(1-5\alpha)(1-\alpha)(t-2)} \qquad (t - 1 \geq s = 1/\alpha) \tag{A.2}$$
$$\geq \lambda \qquad \text{(From the choice of minimal } t\text{)} \tag{A.3}$$

Since $b_2 = 5s \log_2(d_2) = 20s^2 \log_2(s) \leq s^4$ (recall that $s = 1/\alpha \geq 2^{10}$),

$$d_2^{1-2\alpha} = \frac{d_2}{d_2^{2\alpha}} = \frac{d_2}{s^8} \leq \frac{d_2}{b_2^2} = \frac{1}{\lambda_2}.$$

We obtain claim (ii) by the following computation,

$$(d_2)^{(t-1)} \leq \lambda_2^{\frac{-(t-1)}{1-2\alpha}}$$
$$\leq \lambda^{\frac{-2}{(1-2\alpha)(1-5\alpha)(1-\alpha)^2}} \qquad \text{(Using Eq. (A.3))}$$
$$\leq \lambda^{-2(1+10\alpha)}. \qquad \square$$

Lemma A.1.4. *The number of walks of length* $t-1$ *on the* s*-wide replacement product of* X *and* Y *is* $O(|S|/\lambda^{2+\beta})$.

Proof. Since each step of the walk has d_2 options, the number of walks is

$$\begin{aligned}|V(X)||V(Y)| \cdot d_2^{(t-1)} &= n' \cdot d_1^s \cdot d_2^{(t-1)} = n' \cdot d_2^{(t-1)+5s} \\ &= \Theta\left(|S| \cdot d_2^{(t-1)+5s+5}\right) \\ &= O\left(|S| \cdot d_2^{(1+5\alpha)(t-1)}\right).\end{aligned}$$

Using Claim A.1.3 (ii), this implies a size of

$$O\left(|S| \cdot d_2^{(1+5\alpha)(t-1)}\right) = O\left(\frac{|S|}{\lambda^{2(1+10\alpha)(1+5\alpha)}}\right) = O\left(\frac{|S|}{\lambda^{2+32\alpha}}\right) = O\left(\frac{|S|}{\lambda^{2+\beta}}\right). \square$$

Before we prove the main result, we need the following simple observation that will be used to construct a modified $(\varepsilon + o(1))$-biased set starting from an ε-biased set, S. This is needed because the graph obtained from Corollary 4.2.11 does not have size exactly $|S|$ but is only guaranteed to be at most $(1 + o(1))|S|$.

Lemma A.1.5. *Let* S *be an* ε*-biased set of a group* G*. And let* S' *be obtained by adding* $\theta|S|$ *many identity elements. Then,* S' *is an* $(\varepsilon + \theta)$*-biased set.*

Proof. Denote by e the identity element of G. Let ρ be any non-trivial irreducible representation of a group G. From the computation we have,

$$\begin{aligned}\|\mathbb{E}_{s\in S'}\rho(s)\|_{op} &= \frac{1}{1+\theta}\left\|\mathbb{E}_{s\in S}\rho(s) + \theta \cdot \mathbb{E}_{s\in S\setminus S'}\rho(1)\right\|_{op} \\ &\le \|\mathbb{E}_{s\in S}\rho(s)\|_{op} + \theta && (\|\rho(1)\|_{op} = 1) \\ &\le \varepsilon + \theta && (S \text{ is } \varepsilon\text{- biased}). \quad \square\end{aligned}$$

Theorem A.1.6 (Almost Ramanujan Expanders I)**.** *Let* $\mathrm{Cay}(G, S)$ *be* λ_0*-expander with constant* $\lambda_0 \in (0, 1)$*. For every function* $\beta(\lambda) > 0$*, and for any* $\lambda > 0$*, sufficiently small such that*

$$\frac{32}{\beta(\lambda)} \leq \left(\frac{\log(1/\lambda)}{4\log\log(1/\lambda)}\right)^{1/3},$$

there exists a deterministic polynomial time algorithm to construct S' *such that* $\mathrm{Cay}(G, S')$ *is a* λ*-expander with degree* $|S'| = O_{\lambda_0}(|S|/\lambda^{2+\beta})$*. Furthermore, each element in* S' *is the product of* $O(\log(1/\lambda))$ *elements of* S*.*

Proof. We can assume that $s \geq 2^{10}$ since otherwise λ is a constant, and we can just use Theorem 4.2.2.

Initial Boost We first boost the expansion from λ_0 to $1/d_2 \leq \lambda_2^2/3$. Using Theorem 4.2.2 (with its parameter β equal to 1), we can find a new set of generators, S_1, such that $\mathrm{Cay}(G, S_1)$ is $1/d_2$-spectral expander and $|S_1| = O(|S|d_2^5)$. Moreover, we also know that each element in S_1 is a multiple of at most $\log\left(d_2^5\right)$ elements in S. We add multiple copies of the entire set to make the size $|S|d_2^5$.

The s**-wide walk** Obtain an (n', d_1, λ_1) Cayley graph X from Corollary 4.2.11 as explained before. We add $n' - n = O(n^{8/9})$ copies of the identity to S_1 to obtain S_2. By Lemma A.1.5 and the assumption that $s \geq 2^{10}$, S_2 is a $\lambda_2^2/3 + O(n^{-1/9}) \leq 2\lambda_2^2/3$-biased set. We denote by S' the final set of generators obtained by t steps of the s-wide replacement product of X and Y. By definition, each element in S' is a product of t elements in S_2 which has the same elements as S_1. Thus, each element in S' is a product of at most

$$\begin{aligned} O(t\log(d_2)) &\leq O((1+10\alpha)\log(1/\lambda)) && \text{(Using Claim A.1.3 [ii])} \\ &\leq O(\log(1/\lambda)) && \text{(By the assumption that } \alpha \leq 1/128) \end{aligned}$$

elements of S. The only thing that remains is to prove expansion of $\mathrm{Cay}(G, S')$. We pick any irreducible representation ρ and apply Theorem 4.3.13 to the function ρ on $S_2 \leftrightarrow V(X)$. The condition that $2\lambda(X) + \left\|\mathbb{E}_{g\sim S_2}[\rho(g)]\right\|_{\mathrm{op}} \leq \lambda(Y)^2$ translates to $\lambda_1 \leq \lambda_2^2/6$ which is satisfied

by our choice of λ_1. Thus, the final expansion is given by,

$$\begin{aligned}
\left\| \mathop{\mathbb{E}}_{g\in S'}[\rho(g)] \right\|_{op} &:= \left(\lambda_2^s + s\cdot\lambda_2^{s-1} + s^2\cdot\lambda_2^{s-3}\right)^{\lfloor (t-1)/s\rfloor} \\
&\le \left(3s^2\lambda_2^{s-3}\right)^{((t-1)/s)-1} && \left(\text{Using } \lambda_2 \le \frac{20s^2\log s}{s^{2s^2}} \le \frac{1}{3s^2}\right) \\
&\le \left(\lambda_2^{s-4}\right)^{(t-1-s)/s} \\
&\le \lambda_2^{(1-5/s)(1-s/(t-1))(t-1)} \\
&\le \lambda_2^{(1-5\alpha)(1-\alpha)(t-1)} && \left(\text{Using Claim A.1.3 [i]}\right) \\
&= \lambda_2^{(1-5\alpha)(1-\alpha)(t-1)} \le \lambda, && \text{(From the choice of t)}. \qquad \square
\end{aligned}$$

A.2 Signed Non-backtracking Operator

A.2.1 Diagonalizing Non-backtracking Operator

Let $\rho : G \to \mathrm{GL}(\mathbb{C}[G])$ be the regular representation of the group G. More concretely, given an element $g \in G$, the map $\rho(g)e_h = e_{h\cdot g^{-1}}$ where $\mathbb{C}[G] = \mathrm{span}\{e_g \mid g \in G\}$. As an example, let $G = \mathbb{Z}_\ell$, $\mathbb{C}[G] = \mathbb{C}^\ell$. Let P be the $\ell\times\ell$ permutation matrix that maps $Pe_i = e_{i+1}$ where $i+1$ is taken modulo ℓ. Then, $\rho(t) = P^t$ for $t \in \mathbb{Z}_\ell$.

For a map ρ as above and an extended signing s , define a generalized non-backtracking walk matrix in which for a non-zero entry indexed by (e_1, e_2), we replace 1 by the block matrix $\rho(s(e_2))$.

Lemma A.2.1. *The non-backtracking walk matrix of the lifted graph is* $B_{X(s)} = B_X(\rho)$.

Proof. In the lifted graph, the edges are of the form $[(u, i - s(u,v)), (v, i)] =: [u, v, i]$ and thus can be indexed by $E' \times [l]$. The non-backtracking walk matrix $B_{\hat{X}}$ would then have a non zero entry from $([u, v, i], [x, y, j])$ iff $(v, i) = (x, j - s(x, y))$ and $(y, j) \neq (u, i - s(u, v))$. Assume that the first condition is met, i.e., $x = v$ and $j = i + s(x, y)$. If $y = u$, then $i - s(u, v) = i - s(y, x) = i + s(x, y) = j$, and hence, the second condition cannot be met.

This is just a longer way of saying that the lifts give a matching between $u \times G$ and $v \times G$. The implication of all this is that y has to be distinct from u, and thus, the pair of edges $(u, v), (v, y)$ has a non-zero entry in B_χ. Moreover, for every i and every pair of edges $(u, v), (v, y)$, we have a non-zero entry for $(u, v, i), (v, y, i + s(v, y))$ in $B_{\chi(s)}$. Thus, it can be written as a block matrix with the entry in $(u, v), (v, y)$ equal to $\rho(s(v, y))$. □

Since the base graph X and the signing s will be fixed throughout, we will drop the subscript to make reading less hurtful. We will need fact that $\rho(h) = F \operatorname{diag}(\chi_1(h), \cdots, \chi_l(h))F^{-1}$, where χ_i are characters of G and F is a unitary. This follows from Theorem 1.3.8, but can also be derived from the well-known fact that a collection of commuting matrices that are diagonalizable are also simultaneously diagonalizable[3] (see [Con, Thm. 5] for a proof).

Corollary A.2.2. *The non-backtracking walk matrix,* $B_{\chi(s)} = Q \operatorname{diag}(B(\chi_1), \cdots, B(\chi_t))Q^{-1}$ *and thus* $\operatorname{Spec}(B_{\chi(s)}) = \cup_{\chi \in \hat{G}} \operatorname{Spec}(B(\chi))$.

Proof. To ease notation we write $B_\chi(\rho) = \sum M_{u,v} \otimes \rho(s(u, v))$ for some $M_{u,v}$. We have $\rho(s(u, v)) = F \operatorname{diag}(\chi_1(h), \cdots, \chi_l(h))F^{-1}$ and thus,

$$B_\chi(\rho) \;=\; (I \otimes F) \sum M_{u,v} \otimes \operatorname{diag}(\chi_1(h), \cdots, \chi_l(h))(I \otimes F^{-1}) \,.$$

Let $|E| = N$ and let T denote the permutation on Nt that maps $T(i) := bt + (a + 1)$ where a, b are the unique non-negative integers such that $0 \le b < N$ $i - 1 = aN + b$. It can then be seen that $\sum M_{u,v} \otimes \operatorname{diag}(\chi_1(h), \cdots, \chi_t(h)) = T \operatorname{diag}\left(\sum M_{u,v} \otimes \chi_i(h)\right) T^{-1}$. Notice that $\sum M_{u,v} \otimes \chi_i(h) = B_\chi(\chi_i)$ and thus putting it together we have that for $Q = (I \otimes F)T$, $B_{\chi(s)} = Q \operatorname{diag}(B(\chi_1), \cdots, B(\chi_t))Q^{-1}$. The statement on the spectrum follows immediately. □

A.2.2 A Simple Consequence of Ihara-Bass

We now slightly extend a claim in [MOP20] for general signings.

3. Using this, the argument is as follows. Since G is abelian, we have that $\{\rho(h)\}$ are commuting, and since they are invertible, they are diagonalizable. Thus, they simultaneously diagonalize.

Claim A.2.3. *Let* A *be the (signed) adjacency matrix of a* d*-regular graph. Suppose* f *is an eigenvector of* A *satisfying*

$$Af = \left(\beta + \frac{d-1}{\beta}\right) f.$$

Then $g(u,v) \coloneqq (f(u) - \beta f(v))$ *(or in the signed case* $g(u,v) \coloneqq A(u,v)^{-1}(f(u) - \beta \cdot A(u,v)f(v))$*) is an eigenvector of the (signed) non-backtracking matrix* B *with eigenvalue* β*.*

Proof. Let f and g be as in the statement of the claim. Suppose that A and B are not signed. By a direct computation, we have,

$$\begin{aligned}
(Bg)(u,v) &= \sum_{w \sim v, w \neq u} f(v) - \beta \cdot f(w) \\
&= (d-1)f(v) - \sum_{w \sim v, w \neq u} \beta \cdot f(w) \\
&= (d-1)f(v) + \beta \cdot f(u) - \beta \sum_{w \sim v} f(w) \\
&= (d-1)f(v) + \beta \cdot f(u) - \beta(Af)(v) \\
&= (d-1)f(v) + \beta \cdot f(u) - \beta \left(\beta + \frac{d-1}{\beta}\right) f(v) \\
&= \beta(f(u) - \beta \cdot f(v)) = \beta \cdot g(u,v).
\end{aligned}$$

Now suppose that A and B are signed. First note that g is well-defined since for every entry $g(u,v)$ the pair (u,v) is an orientation of an edge of the graph, so it has a signing $A(u,v) \neq 0$. We have

$$\begin{aligned}
(Bg)(u,v) &= \sum_{w \sim v, w \neq u} A(v,w)A(v,w)^{-1}(f(v) - \beta \cdot A(v,w)f(w)) \\
&= (d-1)f(v) - \beta \sum_{w \sim v, w \neq u} A(v,w)f(w) \\
&= (d-1)f(v) + \beta \cdot A(v,u)f(u) - \beta \sum_{w \sim v} A(v,w)f(w)
\end{aligned}$$

$$
\begin{aligned}
&= (d-1)f(v) + \beta \cdot A(v,u)f(u) - \beta(Af)(v) \\
&= (d-1)f(v) + \beta \cdot A(v,u)f(u) - \beta\left(\beta + \frac{d-1}{\beta}\right) f(v) \\
&= \beta \cdot A(v,u)\left(f(u) - \beta\frac{1}{A(v,u)} f(v)\right), \\
&= \beta \cdot A(u,v)^{-1}\left(f(u) - \beta \cdot A(u,v)f(v)\right) = \beta \cdot g(u,v),
\end{aligned}
$$

where we used $A(v,u) = A(u,v)^{-1}$. □

Corollary A.2.4. *Let A be the (signed) adjacency matrix of a d-regular graph. Let B be its (signed) non-backtracking operator. For any $\lambda > 2\sqrt{d-1}$, if $\rho_2(B) \le \lambda/2$, then $\rho_2(A) \le \lambda$. Hence, $\lambda(G) = \rho(A) \le 2\max\left\{\sqrt{d-1}, \rho_2(B)\right\}$.*

Proof. We show via the contrapositive. Suppose that f is eigenvector of A with eigenvalue α such that $|\alpha| > \lambda$. By possibly multiplying A and B by a phase (i.e., $e^{i\theta}$), we can assume α is a non-negative real number. By Claim A.2.3, we have that β satisfying the equation $\beta^2 - \alpha\beta + (d-1) = 0$ is an eigenvalue of B. Considering the solution $\beta^+ = (\alpha + \sqrt{\alpha^2 - 4(d-1)})/2$ and thus, we have $\beta^+ \ge \alpha/2 > \lambda/2$. □

A.3 A Precise Implementation of DFS

We now present the precise implementation of the depth-first search algorithm as we need this implementation to satisfy the following:

Observation A.3.1. Let X be a connected graph. The **DFS** algorithm traversals each edge of X exactly twice: first in a recursive step and subsequently in a backtrack step.

Algorithm A.3.2 (DFS(X, v))**.**

Input connected graph X and starting vertex v

Output traversal $\mathcal{T}$ of G and step types σ

- Color all vertices of X with 'green'
- Traversal $\mathcal{T} = ()$
- Step types $\sigma = ''$
- Parent π = null
- **DFSRec**($X, \mathcal{T}, \pi, \sigma, v, e$ = null)
- return $\mathcal{T}, \sigma$

Algorithm A.3.3 (DFSRec($G, \mathcal{T}, \pi, \sigma, v, e$))**.**

Input graph G, traversal $\mathcal{T}$, parent π, step types σ, vertex v and traversed edge e

Output Updated $\mathcal{T}$ (as side effect)

- $\mathcal{T}$.append(e) if $e \neq$ null
- If v is 'green':
  - Color v with 'yellow'
  - For each neighbor u of v not colored 'red' and $u \neq \pi$:
    - σ.append('R') (recursive step)
    - **DFSRec**($G, \mathcal{T}, \pi = v, \sigma, u, e = (v, u)$)
  - Color v with 'red'
- σ.append('B') (backtrack step)